arXiv:1704.00642v1 [math.ST] 3 Apr 2017

Local nearest neighbour classification with applications to semi-supervised learning

Timothy I. Cannings, Thomas B. Berrett and Richard J. Samworth
Statistical Laboratory, University of Cambridge

Abstract

We derive a new asymptotic expansion for the global excess risk of a local k-nearest neighbour classifier, where the choice of k may depend upon the test point. This expansion elucidates conditions under which the dominant contribution to the excess risk comes from the locus of points at which each class label is equally likely to occur, but we also show that if these conditions are not satisfied, the dominant contribution may arise from the tails of the marginal distribution of the features. Moreover, we prove that, provided the d-dimensional marginal distribution of the features has a finite ρth moment for some $\rho > 4$ (as well as other regularity conditions), a local choice of k can yield a rate of convergence of the excess risk of $O(n^{-4/(d+4)})$, where n is the sample size, whereas for the standard k-nearest neighbour classifier, our theory would require $d \geq 5$ and $\rho > 4d/(d-4)$ finite moments to achieve this rate. Our results motivate a new k-nearest neighbour classifier for semi-supervised learning problems, where the unlabelled data are used to obtain an estimate of the marginal feature density, and fewer neighbours are used for classification when this density estimate is small. The potential improvements over the standard k-nearest neighbour classifier are illustrated both through our theory and via a simulation study.

1 Introduction

Supervised classification problems represent some of the most frequently-occurring statistical challenges in a wide variety of fields, including fraud detection, medical diagnoses and targeted advertising, to name just a few. The area has received an enormous amount of attention within both the statistics and machine learning communities; for an excellent survey with pointers to much of the relevant literature, see Boucheron, Bousquet and Lugosi (2005).

The k-nearest neighbour classifier, which assigns the test point according to a majority vote over the classes of its k nearest points in the training set, is arguably the simplest and most intuitive nonparametric classifier. It was introduced in the seminal work of Fix and Hodges (1951), later republished as Fix and Hodges (1989), and early understanding of some of its theoretical properties was provided in Cover and Hart (1967), Duda and Hart (1973) and Stone (1977). Further recent contributions, some of which treat the k-nearest neighbour classifier as a special case of a plug-in classifier, include Kulkarni and Posner (1995), Audibert and Tsybakov (2007), Hall et al. (2008), Biau et al. (2010), Samworth (2012), Chaudhuri and Dasgupta (2014) and Celisse and Mary-Huard (2015). Nearest neighbour methods have also been extensively used in other statistical problems, including nonparametric clustering (Heckel and Bölcskei, 2015), entropy estimation (Kozachenko and Leonenko, 1987; Berrett et al., 2016) and testing problems (Schilling, 1986); see also the recent book Biau and Devroye (2015).

Despite these aforementioned works, the behaviour of the k-nearest neighbour classifier in the tails of a distribution remains poorly understood. Indeed, writing (X, Y) for a generic data pair, where the d-dimensional feature vector X has marginal density $\bar{f}$ and Y denotes a binary class label, most of the results in the papers mentioned in the previous paragraph pertain either to situations where $\bar{f}$ is compactly supported and bounded away from zero on its support, or where the excess risk is computed only over a compact subset of $\mathbb{R}^d$. Unfortunately, such restrictions are typically imposed purely for mathematical convenience, and leave open the question of the effect of tail behaviour on the excess risk.

The first goal of this paper, therefore, is to provide a new asymptotic expansion for the global excess risk of a k-nearest neighbour classifier (Theorem 1), where we allow the feature vectors to have unbounded support. Our expansion elucidates conditions under which the dominant contribution to the excess risk comes from the locus of points at which each class label is equally likely to occur, but we also show that if these conditions are not satisfied, the dominant contribution may arise from the tails of the marginal distribution of the features.

The proof of Theorem 1 also reveals a local bias-variance trade-off that motivates a modification of the k-nearest neighbour classifier in semi-supervised classification settings, where, in addition to the labelled training data, we have access to a further, independent, sample of unlabelled observations. Such semi-supervised problems occur in a wide range of applications, especially where it is expensive or time-consuming to obtain the labels associated with observations; in fact, it is frequently the case that unlabelled observations may vastly outnumber labelled ones. For an overview of semi-supervised learning applications and techniques, see Chapelle et al. (2006).

Our second contribution is to propose to allow the choice of k to depend on an estimate of $\bar{f}$ at the test point in semi-supervised settings. By using fewer neighbours in low density regions, we are able to achieve a better balance in the local bias-variance trade-off. In particular, we initially study an oracle, local choice of k that depends on $\bar{f}$, and under regularity conditions, we show that the excess risk over $\mathbb{R}^d$ is $O(n^{-4/(d+4)})$ provided that the feature vectors have $\rho > 4$ finite moments. By contrast, our theory for the standard k-nearest neighbour classifier with a global choice of k requires that $d \geq 5$ and the feature vectors have $\rho > 4d/(d-4)$ finite moments. Assuming further that $\bar{f}$ has Hölder smoothness $\gamma \in (0, 2]$, we show that if m additional, unlabelled observations are used to estimate $\bar{f}$ by $\hat{f}_m$, and if $m = m_n$ satisfies $\liminf_{n\to\infty} m_n/n^{2+d/\gamma} > 0$, then our semi-supervised k-nearest-neighbour classifier mimics the asymptotic performance of the oracle.

As mentioned previously, studies of global excess risk rates of convergence in nonparametric classification for unbounded feature vector distributions are comparatively rare. Hall and Kang (2005) studied the tail error properties of a classifier based on kernel density estimates of the class conditional densities for univariate data. As an illustrative example, they showed that if, for large x, one class has density $ax^{-\alpha}$, while the other has density $bx^{-\beta}$, for some $a, b > 0$ and $1 < \alpha < \beta < \alpha + 1 < \infty$, then the excess risk from the right tail is of larger order than that in the body of the distribution.

Perhaps most closely related to this work, Gadat et al. (2016) recently obtained upper bounds on the supremum excess risk of the k-nearest neighbour classifier, when η is Lipschitz, the well-known margin assumption of Mammen and Tsybakov (1999) is satisfied, and a tail condition on the rate of decay of $\mathbb{P}\{\bar{f}(X) < \delta\}$ as $\delta \searrow 0$ is imposed. They also derived minimax lower bounds (in general, of different order) in the same problem. Our assumptions and conclusions are not directly comparable, but allow us to obtain the same rates of convergence as in situations where the marginal distribution of X is compactly supported and bounded away from zero on its support, as well as to provide the leading constants in the asymptotic expansion for the excess risk in such cases.

The remainder of this paper is organised as follows. After introducing our setting in Section 2, we present in Section 3 our main results for the standard k-nearest neighbour classifier. This leads on, in Section 4, to our study of the semi-supervised setting, where we derive asymptotic results of the excess risk of our local k-nearest neighbour classifier. We illustrate the finite-sample benefits of the semi-supervised classifier over the standard k-nearest neighbour classifier in a simulation study in Section 5. Proofs are given in Section 6, while in the Appendix, we present an introduction to the ideas of differential geometry that underpin much of our analysis.

Finally we fix here some notation used throughout the paper. Let $\|\cdot\|$ denote the Euclidean norm and, for $r > 0$ and $x \in \mathbb{R}^d$, let $B_r(x) := \{z \in \mathbb{R}^d : \|x - z\| < r\}$ and $\bar{B}_r(x) := \{z \in \mathbb{R}^d : \|x - z\| \leq r\}$ denote respectively the open and closed Euclidean balls of radius r centred at x. Let $a_d := \frac{2\pi^{d/2}}{d\Gamma(d/2)}$ denote the d-dimensional Lebesgue measure of $B_1(0)$. For a real-valued function g defined on $A \subseteq \mathbb{R}^d$ that is twice differentiable at x, write $\dot{g}(x) = (g_1(x), \ldots, g_d(x))^T$ and $\ddot{g}(x) = \big(g_{jk}(x)\big)$ for its gradient vector and Hessian matrix at x, and let $\|g\|_\infty = \sup_{x \in A} |g(x)|$. Let $\|\cdot\|_{\mathrm{op}}$ denote the operator norm of a matrix.

2 Statistical setting

Let $(X, Y), (X_1, Y_1), \ldots, (X_{n+m}, Y_{n+m})$ be independent and identically distributed random pairs taking values in $\mathbb{R}^d \times \{1, 2\}$. Let $\pi_r := \mathbb{P}(Y = r)$, for $r = 1, 2$, and $X|Y = r \sim P_r$, for $r = 1, 2$, where P_r is a probability measure on $\mathbb{R}^d$. Let $\eta(x) := \mathbb{P}(Y = 1|X = x)$ and let $P_X := \pi_1 P_1 + \pi_2 P_2$ denote the marginal distribution of X. We observe *labelled training data*, $\mathcal{T}_n := \{(X_1, Y_1), \ldots, (X_n, Y_n)\}$, and *unlabelled training data*, $\mathcal{T}'_m := \{X_{n+1}, \ldots, X_{n+m}\}$, and are presented with the task of assigning the *test point* X to either class 1 or 2.

A *classifier* is a Borel measurable function $C : \mathbb{R}^d \to \{1, 2\}$, with the interpretation that C assigns $x \in \mathbb{R}^d$ to the class $C(x)$. Given a Borel measurable set $\mathcal{R} \subseteq \mathbb{R}^d$, the misclassification rate, or *risk*, over $\mathcal{R}$ is

$$R_{\mathcal{R}}(C) := \mathbb{P}[\{C(X) \neq Y\} \cap \{X \in \mathcal{R}\}].$$

When $\mathcal{R} = \mathbb{R}^d$, we drop the subscript for convenience. The *Bayes classifier*

$$C^{\mathrm{Bayes}}(x) := \begin{cases} 1 & \text{if } \eta(x) \geq 1/2; \\ 2 & \text{otherwise,} \end{cases}$$

minimises the risk over any region $\mathcal{R}$ (Devroye et al., 1996, p. 20). Thus, the performance of a classifier C is measured via its (non-negative) *excess risk*, $R_{\mathcal{R}}(C) - R_{\mathcal{R}}(C^{\mathrm{Bayes}})$.

We can now formally define the local-k-nearest neighbour classifier, which allows the number of neighbours considered to vary depending on the location of the test point. Suppose $k_{\mathrm{L}} : \mathbb{R}^d \to \{1, \ldots, n\}$ is measurable. Given the test point $x \in \mathbb{R}^d$, let $(X_{(1)}, Y_{(1)}), \ldots, (X_{(n)}, Y_{(n)})$ be a reordering of the training data such that $\|X_{(1)} - x\| \leq \cdots \leq \|X_{(n)} - x\|$. We will later assume that P_X is absolutely continuous with respect to d-dimensional Lebesgue measure, which ensures that ties occur with probability zero; where helpful for clarity, we also write $X_{(i)}(x)$ for the ith nearest neighbour of x. Let $\hat{S}_n(x) := k_{\mathrm{L}}(x)^{-1} \sum_{i=1}^{k_{\mathrm{L}}(x)} \mathbb{1}_{\{Y_{(i)}=1\}}$. Then the

local-k-nearest neighbour (k_{L}nn) classifier is defined to be

$$\hat{C}_n^{k_{\mathrm{L}}\mathrm{nn}}(x) := \begin{cases} 1 & \text{if } \hat{S}_n(x) \geq 1/2; \\ 2 & \text{otherwise.} \end{cases}$$

Given $k \in \{1, \ldots, n\}$, let k_0 denote the constant function $k_0(x) := k$ for all $x \in \mathbb{R}^d$. Using $k_{\mathrm{L}} = k_0$ the definition above reduces to the standard *k-nearest neighbour classifier* ($k$nn), and we will write $\hat{C}_n^{k\mathrm{nn}}$ in place of $\hat{C}_n^{k_0\mathrm{nn}}$. For $\beta \in (0, 1/2)$, let

$$K_\beta := \{\lceil (n-1)^\beta \rceil, \lceil (n-1)^\beta \rceil + 1, \ldots, \lfloor (n-1)^{1-\beta} \rfloor\}$$

denote a range of values of k that will be of interest to us. Note that $K_{\beta_1} \supseteq K_{\beta_2}$, for $\beta_1 < \beta_2$. Moreover, when β is small, the upper and lower bounds are only slightly stronger requirement than the consistency conditions of Stone (1977), namely that $k = k_n \to \infty$, $k_n/n \to 0$ as $n \to \infty$.

3 Global risk of the k-nearest neighbour classifier

Our aim in this section is to provide an asymptotic expansion for the global risk of the standard (non-local) k-nearest neighbour classifier. Our analysis will make use of the following assumptions:

(A.1) The probability measures P_1 and P_2 are absolutely continuous with respect to Lebesgue measure, with Radon–Nikodym derivatives f_1 and f_2, respectively. Moreover, the marginal density of X, given by $\bar{f} := \pi_1 f_1 + \pi_2 f_2$, is continuous P_X-almost everywhere and $\mathcal{X}_{\bar{f}} := \{x \in \mathbb{R}^d : \bar{f} \text{ is continuous at } x\}$ is open.

Let $\mathcal{S} := \{x \in \mathbb{R}^d : \eta(x) = 1/2\}$ and, for $\epsilon > 0$, let $\mathcal{S}^\epsilon := \mathcal{S} + B_\epsilon(0)$.

(A.2) The set $\mathcal{S} \cap \{x \in \mathbb{R}^d : \bar{f}(x) > 0\}$ is non-empty and $\bar{f}$ is bounded on $\mathcal{S}$. There exist $\epsilon_0 > 0$ and a measurable function $g : \mathcal{S} \to [1, \infty)$ with the property that $\bar{f}$ is twice continuously differentiable on $\mathcal{S}^{\epsilon_0}$, and

$$\max\left\{ \|\dot{\bar{f}}(x_0)\|, \sup_{u \in B_{\epsilon_0}(0)} \|\ddot{\bar{f}}(x_0 + u)\|_{\mathrm{op}} \right\} \leq \bar{f}(x_0) g(x_0), \tag{1}$$

for all $x_0 \in \mathcal{S}$, where $\sup_{x_0 \in \mathcal{S} : \bar{f}(x_0) \geq \delta} g(x_0) = o(\delta^{-\tau})$, as $\delta \searrow 0$, for each $\tau > 0$. Furthermore, writing $p_\epsilon(x) := P_X\big(B_\epsilon(x)\big)$, there exists $\mu_0 \in (0, a_d)$ such that, for all $x \in \mathbb{R}^d \setminus \mathcal{S}^{\epsilon_0}$ and $\epsilon \in (0, \epsilon_0]$, we have

$$p_\epsilon(x) \geq \mu_0 \epsilon^d \bar{f}(x).$$

(A.3) We have $\inf_{x_0 \in \mathcal{S}} \|\dot{\eta}(x_0)\| > 0$, so that $\mathcal{S}$ is a $(d-1)$-dimensional, orientable manifold (cf. Section 7.3). Moreover, $\sup_{x \in \mathcal{S}^{2\epsilon_0}} \|\dot{\eta}(x)\| < \infty$ and $\ddot{\eta}$ is uniformly continuous on $\mathcal{S}^{2\epsilon_0}$ with $\sup_{x \in \mathcal{S}^{2\epsilon_0}} \|\ddot{\eta}(x)\|_{\mathrm{op}} < \infty$. Finally, the function η is continuous on $\{x : \bar{f}(x) > 0\}$, and for every $\tau > 0$,

$$\sup_{x \in \mathbb{R}^d \setminus \mathcal{S}^{\epsilon_0} : \bar{f}(x) \geq \delta} |\eta(x) - 1/2|^{-1} = o(\delta^{-\tau})$$

as $\delta \searrow 0$.

(A.4)(ρ) We have that $\int_{\mathbb{R}^d} \|x\|^\rho dP_X(x) < \infty$ and $\int_{\mathcal{S}} \bar{f}(x_0)^{d/(\rho+d)} \, d\mathrm{Vol}^{d-1}(x_0) < \infty$, where $d\mathrm{Vol}^{d-1}$ denotes the $(d-1)$-dimensional volume form on $\mathcal{S}$ (cf. Section 7.3).

The density assumption in **(A.1)** allows us to define the tail of the distribution as the region where $\bar{f}$ is smaller than some threshold. The second and third parts of **(A.1)** ensure that for all $\delta > 0$ sufficiently small, the set $\mathcal{R} := \{x : \bar{f}(x) > \delta\} \cap \mathcal{X}_{\bar{f}}$ is a d-dimensional manifold, and $P_X(\mathcal{R}^c) \leq \mathbb{P}\{\bar{f}(X) \leq \delta\}$, where the latter quantity can be bounded straightforwardly using **(A.4)**(ρ). The first part of **(A.2)** asks for a certain level of smoothness for $\bar{f}$ in a neighbourhood of $\mathcal{S}$, and controls the behaviour of its first and second derivatives there relative to the original density. In particular, the greater degree of regularity asked of these derivatives in the tails of the marginal density allows us still to control the error of a Taylor approximation even in this region. Moreover, (1) is satisfied by all Gaussian and multivariate-t densities, for example. The second part of **(A.2)** concerns the behaviour of the marginal feature distribution away from $\mathcal{S}^{\epsilon_0}$ and is often referred to as the strong minimal mass assumption (e.g. Gadat et al., 2016). It requires that the mass of the marginal feature distribution is not concentrated in the neighbourhood of a point and is a rather weaker condition than we ask for on $\mathcal{S}^{\epsilon_0}$; in particular, we do not ask for derivatives of $\bar{f}$ in this region.

The first part of **(A.3)** asks for the class conditional densities, when weighted by their respective prior probabilities, to cross at an angle, while the bounds on the first and second derivatives of η ensure that we can estimate η sufficiently well. The last part of this condition asks that η does not approach the critical value of $1/2$ too fast on the complement of $\mathcal{S}^{\epsilon_0}$. Finally, the first condition of **(A.4)**(ρ) is a simple moment condition, while the second ensures the constants B_1 and B_2 in (2) below are finite.

Let

$$B_1 := \int_{\mathcal{S}} \frac{\bar{f}(x_0)}{4\|\dot{\eta}(x_0)\|} \, d\mathrm{Vol}^{d-1}(x_0) \text{ and } B_2 := \int_{\mathcal{S}} \frac{\bar{f}(x_0)^{1-4/d}}{\|\dot{\eta}(x_0)\|} a(x_0)^2 \, d\mathrm{Vol}^{d-1}(x_0), \tag{2}$$

where

$$a(x) := \frac{\sum_{j=1}^{d}\{\eta_j(x)\bar{f}_j(x) + \frac{1}{2}\eta_{jj}(x)\bar{f}(x)\}}{(d+2)a_d^{2/d}\bar{f}(x)}. \tag{3}$$

We are now in a position to present our asymptotic expansion for the global excess risk of the standard k-nearest neighbour classifier.

Theorem 1. *Assume* ***(A.1)***, ***(A.2)***, ***(A.3)*** *and* ***(A.4)***(ρ).

(i) Suppose that $d \geq 5$ *and* $\rho > \frac{4d}{d-4}$. *Then for each* $\beta \in (0, 1/2)$,

$$R(\hat{C}_n^{k\mathrm{nn}}) - R(C^{\mathrm{Bayes}}) = \frac{B_1}{k} + B_2\Big(\frac{k}{n}\Big)^{4/d} + o\bigg(\frac{1}{k} + \Big(\frac{k}{n}\Big)^{4/d}\bigg)$$

as $n \to \infty$, *uniformly for* $k \in K_\beta$.

(ii) Suppose that either $d \leq 4$, *or,* $d \geq 5$ *and* $\rho \leq \frac{4d}{d-4}$. *Then for each* $\beta \in (0, 1/2)$ *and each* $\epsilon > 0$ *we have*

$$R(\hat{C}_n^{k\mathrm{nn}}) - R(C^{\mathrm{Bayes}}) = \frac{B_1}{k} + o\bigg(\frac{1}{k} + \Big(\frac{k}{n}\Big)^{\frac{\rho}{\rho+d}-\epsilon}\bigg)$$

as $n \to \infty$, *uniformly for* $k \in K_\beta$.

Theorem 1 reveals an interesting dichotomy: we see from part (i) that, when $d \geq 5$ and **(A.4)**(ρ) holds for sufficiently large ρ (and the other regularity conditions hold), the dominant contribution to the excess risk arises from the difficulty of classifying points close to the Bayes decision boundary $\mathcal{S}$. In such settings, the excess risk of the standard k-nearest neighbour classifier converges to zero at rate $O(n^{-4/(d+4)})$ when k is chosen proportional to $n^{4/(d+4)}$. On the other hand, part (ii) suggests that when either $d \leq 4$ or $d \geq 5$ and we only know that **(A.4)**(ρ) holds for small ρ, the dominant contribution to the excess risk when k is large may come from the challenge of classifying points in the tails of the distribution. Indeed, Example 1 below provides one simple setting where this dominant contribution does come from the tails of the distribution.

The proof of Theorem 1, and indeed those of Theorems 2 and 3 which follow in Section 4 below, depend crucially on Theorem 4 in Section 6. This result provides an asymptotic expansion for the excess risk of a general (local or global) k-nearest neighbour classifier over a region $\mathcal{R}_n \subseteq \{x \in \mathbb{R}^d : \bar{f}(x) \geq \delta_n(x)\}$, where $\delta_n(x)$, defined in (8) below, shrinks to zero at a rate slow enough to ensure that $X_{(k)}(x)$ concentrates around x uniformly over $\mathcal{R}_n$. This enables us to derive asymptotic expansions for the bias and variance of $\hat{S}_n(x)$, uniformly over $\mathcal{R}_n$, and using a normal approximation, we can deduce an asymptotic expansion for the excess risk, uniformly over the relevant set of nearest neighbour classifiers. Having proved Theorem 4, the proof of Theorem 1 is completed by controlling the remainder terms in Theorem 4 appropriately, and bounding $P_X(\mathcal{R}_n^c)$ using **(A.4)**(ρ).

Example 1. *Suppose that the joint density of X at $x = (x_1, x_2) \in (0,1) \times \mathbb{R}$ is given by*

$$\bar{f}(x) = 2x_1 f_2(x_2),$$

*where f_2 is a positive, twice continuously differentiable density with $f_2(x_2) = e^{-|x_2|}/2$ for $|x_2| > 1$. Suppose also that $\eta(x) = x_1$. Then **(A.1)**, **(A.2)**, **(A.3)** hold, and **(A.4)**(ρ) holds for every $\rho > 0$. We prove in Section 6.3 that for every $\beta \in (0, 1/2)$ and $\epsilon > 0$,*

$$\liminf_{n\to\infty} \inf_{k \in K_\beta} \Big(\frac{n}{k}\Big)^{1+\epsilon} \big\{ R(\hat{C}_n^{k\mathrm{nn}}) - R(C^{\mathrm{Bayes}}) \big\} > 0 \tag{4}$$

as $n \to \infty$.

4 Local k-nearest neighbour classifiers

In this section we explore the consequences of a local choice of k, compared with the global choice in Theorem 1. Initially, we consider an oracle choice, where k is allowed to depend on the marginal feature density $\bar{f}$ (Section 4.1), but we then relax this to semi-supervised settings, where $\bar{f}$ can be estimated from unlabelled training data (Section 4.2).

4.1 Oracle classifier

Suppose for now that the marginal density $\bar{f}$ is known. For $\beta \in (0, 1/2)$ and $B > 0$, let

$$k_{\mathrm{O}}(x) := \max\Big[\lceil (n-1)^\beta \rceil \,,\, \min\big\{ \lfloor B\big\{ \bar{f}(x)(n-1) \big\}^{4/(d+4)} \rfloor \,,\, \lfloor (n-1)^{1-\beta} \rfloor \big\} \Big], \tag{5}$$

where the subscript O refers to the fact that this is an oracle choice of the function k_{L}, since it depends on $\bar{f}$. This choice aims to balance the local bias and variance of $\hat{S}_n(x)$.

Theorem 2. *Assume **(A.1)**, **(A.2)**, **(A.3)** and **(A.4)**(ρ). Then for each $0 < B_* \le B^* < \infty$,*

(i) if $\rho > 4$ then for $\beta < 4d(\rho-4)/\{\rho(d+4)^2\}$,

$$R(\hat{C}_n^{k_{\mathrm{O}}\mathrm{nn}}) - R(C^{\mathrm{Bayes}}) = B_3 n^{-4/(d+4)}\{1 + o(1)\},$$

uniformly for $B \in [B_*, B^*]$ as $n \to \infty$, where

$$B_3 := \int_{\mathcal{S}} \frac{\bar{f}(x_0)^{d/(d+4)}}{\|\dot{\eta}(x_0)\|} \Big\{ \frac{1}{4B} + B^{4/d} a(x_0)^2 \Big\} \, d\mathrm{Vol}^{d-1}(x_0).$$

(ii) if $\rho \le 4$ and $\beta < \min\{1/2, 4/(d+4)\}$, then for every $\epsilon > 0$

$$R(\hat{C}_n^{k_{\mathrm{O}}\mathrm{nn}}) - R(C^{\mathrm{Bayes}}) = o(n^{-\rho/(\rho+d)+\beta+\epsilon}),$$

uniformly for $B \in [B_*, B^*]$, as $n \to \infty$.

Comparing Theorem 2(i) and Theorem 1(i), we see that, unlike for the global k-nearest neighbour classifier, we can guarantee a $O(n^{-4/(d+4)})$ rate of convergence for the excess risk of the oracle classifier, both in low dimensions ($d \leq 4$), and under a weaker condition on ρ when $d \geq 5$. In particular, the condition on ρ no longer depends on the dimension of the covariates. The guarantees in Theorem 2(ii) are also stronger than those provided by Theorem 1(ii) for any global choice of k. Examining the proof of Theorem 2, we find that the key difference with the proof of Theorem 1 is that we can now choose the region $\mathcal{R}_n$ (cf. the discussion following the statement of Theorem 1) to be larger.

4.2 The semi-supervised nearest neighbour classifier

Now consider the more realistic setting where the marginal density $\bar{f}$ of X is unknown, but where we have access to an estimate $\hat{f}_m$ based on the unlabelled training set $\mathcal{T}'_m$. Of course, many different techniques are available, but for simplicity, we focus here on a kernel method. Let K be a bounded kernel with $\int_{\mathbb{R}^d} K(x)\,dx = 1$, $\int_{\mathbb{R}^d} xK(x)\,dx = 0$, $\int_{\mathbb{R}^d} \|x\|^2|K(x)|\,dx < \infty$, and let $R(K) := \int_{\mathbb{R}^d} K(x)^2\,dx$. We further assume that $K(x) = Q(p(x))$, where p is a polynomial and Q is a function of bounded variation. Now define a kernel density estimator of $\bar{f}$, given by

$$\hat{f}_m(x) = \hat{f}_{m,h}(x) := \frac{1}{mh^d} \sum_{j=1}^{m} K\Big(\frac{x - X_{n+j}}{h}\Big).$$

Motivated by the oracle local choice of k in (5), for $\beta \in (0, 1/2)$ and $B > 0$, let

$$k_{\mathrm{SS}}(x) := \max\Big[\lceil (n-1)^{\beta} \rceil\,,\, \min\big\{\lfloor B\{\hat{f}_m(x)(n-1)\}^{4/(d+4)} \rfloor\,,\, \lfloor (n-1)^{1-\beta} \rfloor\big\}\Big].$$

For $\gamma \in (0, 2]$, we will consider the following condition:

(A.5)(γ) We have that $\bar{f}$ is bounded and, if $\gamma > 1$, then $\bar{f}$ is differentiable on $\mathbb{R}^d$; moreover, there exists $\lambda > 0$ such that

$$\|\bar{f}(y) - \bar{f}(x)\| \leq \lambda \|y - x\|^{\gamma} \quad \text{for all } x, y \in \mathbb{R}^d, \text{ if } \gamma \in (0, 1]$$
$$\|\dot{\bar{f}}(y) - \dot{\bar{f}}(x)\| \leq \lambda \|y - x\|^{\gamma - 1} \text{ for all } x, y \in \mathbb{R}^d, \text{ if } \gamma \in (1, 2].$$

Theorem 3. *Assume **(A.1)**, **(A.2)**, **(A.3)**, **(A.4)(ρ)** and **(A.5)(γ)** for some $\gamma \in (0, 2]$. Let $m_0 > 0$, let $0 < A_* \leq A^* < \infty$ and $0 < B_* \leq B^* < \infty$, and let $h = h_m := Am^{-1/(d+2\gamma)}$ for some $A > 0$.*

(i) If $\rho > 4$ and $\beta < 4d(\rho - 4)/\{\rho(d+4)^2\}$,

$$R(\hat{C}_n^{k_{\mathrm{SS}}\mathrm{nn}}) - R(C^{\mathrm{Bayes}}) = B_3 n^{-4/(d+4)}\{1 + o(1)\}$$

uniformly for $A \in [A_*, A^*]$, $B \in [B_*, B^*]$ *and* $m = m_n \geq m_0(n-1)^{2+d/\gamma}$, *where* B_3 *was defined in Theorem 2(i).*

(ii) if $\rho \leq 4$ *and* $\beta < \min\{1/2, 4/(d+4)\}$, *then for every* $\epsilon > 0$,

$$R(\hat{C}_n^{k_{\mathrm{SS}}\mathrm{nn}}) - R(C^{\mathrm{Bayes}}) = o(n^{-\rho/(\rho+d)+\beta+\epsilon}),$$

uniformly for $A \in [A_*, A^*]$, $B \in [B_*, B^*]$ *and* $m = m_n \geq m_0(n-1)^{2+d/\gamma}$.

Examination of the proof of Theorem 3 reveals that the same conclusion holds for any estimator $\tilde{f}_m$ of $\bar{f}$ constructed from $\mathcal{T}'_m$, provided there exists $\alpha > (1+d/4)\beta$ such that

$$\mathbb{P}\left(\|\tilde{f}_m - \bar{f}\|_\infty \geq \frac{1}{(n-1)^{1-\alpha/2}} \right) = o(n^{-4/(d+4)}). \tag{6}$$

Condition **(A.5)**(γ) ensures that (6) holds for our kernel density estimator.

5 Empirical analysis

In this section, we compare the k_{O}nn and k_{SS}nn classifiers, introduced in Section 4 above, with the standard $k$nn classifier studied in Section 3. We investigate three settings that reflect the differences between the main results in Sections 3 and 4.

- Setting 1: P_1 is the distribution of d independent $N(0,1)$ components; whereas P_2 is the distribution of d independent $N(1,1/4)$ components.
- Setting 2: P_1 is the distribution of d independent t_5 components; P_2 is the distribution of d independent components, the first $\lfloor d/2 \rfloor$ having a t_5 distribution and the remaining $d - \lfloor d/2 \rfloor$ having a $N(1,1)$ distribution.
- Setting 3: P_1 is the distribution of d independent standard Cauchy components; P_2 is the distribution of d independent components, the first $\lfloor d/2 \rfloor$ being standard Cauchy and the remaining $d - \lfloor d/2 \rfloor$ standard normal.

The corresponding marginal distribution P_X in Setting 1 has all moments finite. Hence, for the standard k-nearest neighbour classifier when $d \geq 5$, we are in the setting of Theorem 1(i), while for $d \leq 4$, we can only appeal to Theorem 1(ii). On the other hand, for the local k-nearest neighbour classifiers, the results of Theorems 2(i) and 3(i) apply for all dimensions, and we can expect the excess risk to converge to zero at rate $O(n^{-4/(d+4)})$. In Setting 2, **(A.4)**(ρ) holds for $\rho < 5$, but not for $\rho \geq 5$. Thus, for the standard k-nearest neighbour classifier, we are in the setting of Theorem 1(ii) for $d < 20$, whereas Theorems 2(i)

	d	Bayes risk	n	$\hat{k}$nn risk	$\hat{k}_{\mathrm{O}}$nn risk	$\hat{k}_{\mathrm{SS}}$nn risk	O RR	SS RR
Setting 1	1	22.67	50	$26.85_{0.13}$	$25.91_{0.12}$	$25.98_{0.13}$	$0.78_{0.022}$	$0.79_{0.023}$
			200	$24.07_{0.06}$	$23.52_{0.06}$	$23.48_{0.05}$	$0.61_{0.030}$	$0.58_{0.029}$
			1000	$23.20_{0.04}$	$22.93_{0.04}$	$22.94_{0.04}$	$0.48_{0.048}$	$0.50_{0.048}$
	2	13.30	50	$17.70_{0.09}$	$16.96_{0.08}$	$16.95_{0.08}$	$0.83_{0.015}$	$0.83_{0.015}$
			200	$15.09_{0.05}$	$14.69_{0.04}$	$14.74_{0.05}$	$0.77_{0.018}$	$0.80_{0.019}$
			1000	$14.04_{0.04}$	$13.78_{0.03}$	$13.80_{0.03}$	$0.65_{0.025}$	$0.67_{0.025}$
	5	3.53	50	$9.46_{0.07}$	$8.95_{0.06}$	$8.94_{0.06}$	$0.91_{0.006}$	$0.91_{0.006}$
			200	$6.94_{0.03}$	$6.67_{0.03}$	$6.70_{0.03}$	$0.92_{0.006}$	$0.93_{0.007}$
			1000	$5.49_{0.02}$	$5.18_{0.02}$	$5.23_{0.02}$	$0.84_{0.008}$	$0.87_{0.008}$
Setting 2	1	31.16	50	$36.55_{0.14}$	$36.07_{0.14}$	$35.93_{0.14}$	$0.91_{0.020}$	$0.88_{0.020}$
			200	$32.93_{0.08}$	$32.38_{0.07}$	$32.42_{0.07}$	$0.69_{0.031}$	$0.71_{0.032}$
			1000	$31.62_{0.05}$	$31.37_{0.05}$	$31.37_{0.05}$	$0.46_{0.065}$	$0.47_{0.066}$
	2	31.15	50	$37.79_{0.13}$	$38.02_{0.12}$	$37.90_{0.12}$	$1.02_{0.014}$	$1.01_{0.015}$
			200	$33.64_{0.08}$	$33.63_{0.07}$	$33.54_{0.07}$	$1.00_{0.028}$	$0.96_{0.026}$
			1000	$31.83_{0.05}$	$31.81_{0.05}$	$31.80_{0.05}$	$0.97_{0.039}$	$0.95_{0.038}$
	5	20.10	50	$28.74_{0.12}$	$29.16_{0.12}$	$29.13_{0.11}$	$1.05_{0.011}$	$1.05_{0.011}$
			200	$23.60_{0.06}$	$23.75_{0.06}$	$23.93_{0.06}$	$1.04_{0.014}$	$1.09_{0.015}$
			1000	$21.86_{0.04}$	$21.71_{0.04}$	$21.77_{0.04}$	$0.91_{0.014}$	$0.95_{0.014}$
Setting 3	1	37.44	50	$44.76_{0.10}$	$43.09_{0.12}$	$43.08_{0.12}$	$0.77_{0.013}$	$0.77_{0.013}$
			200	$41.86_{0.08}$	$40.18_{0.09}$	$40.23_{0.09}$	$0.62_{0.017}$	$0.63_{0.017}$
			1000	$38.68_{0.06}$	$37.85_{0.05}$	$37.89_{0.05}$	$0.33_{0.033}$	$0.36_{0.032}$
	2	37.45	50	$46.20_{0.09}$	$44.81_{0.10}$	$45.24_{0.10}$	$0.84_{0.009}$	$0.89_{0.009}$
			200	$43.50_{0.07}$	$42.29_{0.08}$	$42.86_{0.08}$	$0.80_{0.011}$	$0.89_{0.011}$
			1000	$40.53_{0.06}$	$39.64_{0.06}$	$39.96_{0.06}$	$0.71_{0.013}$	$0.82_{0.014}$
	5	23.23	50	$41.56_{0.11}$	$38.13_{0.11}$	$39.26_{0.12}$	$0.81_{0.005}$	$0.87_{0.005}$
			200	$36.02_{0.07}$	$33.34_{0.06}$	$34.68_{0.07}$	$0.79_{0.004}$	$0.90_{0.004}$
			1000	$31.46_{0.05}$	$29.91_{0.05}$	$30.58_{0.05}$	$0.81_{0.004}$	$0.89_{0.004}$

Table 1: Misclassification rates for Settings 1, 2 and 3. In the final two columns we present the regret ratios given in (7) (with standard errors calculated via the delta method).

and 3(i) again apply for all dimensions for the local classifiers. Finally, in Setting 3, **(A.4)**(ρ) does not hold for any $\rho \geq 1$, and only the conditions of Theorems 1(ii), 2(ii) and 3(ii) apply.

For the standard $k$nn classifier, we use 5-fold cross validation to choose k, based on a

sequence of equally-spaced values between 1 and $\lfloor n/4 \rfloor$ of length at most 40. For the oracle classifier, we set

$$\hat{k}_{\mathrm{O}}(x) := \max\Big[1, \min\big[\lfloor \hat{B}_{\mathrm{O}}\{\bar{f}(x)n/\|\bar{f}\|_\infty\}^{4/(d+4)} \rfloor, n/2\big]\Big],$$

where $\hat{B}_{\mathrm{O}}$ was again chosen via 5-fold cross validation, but based on a sequence of 40 equally-spaced points between $n^{-4/(d+4)}$ (corresponding to the 1-nearest neighbour classifier) and $n^{d/(d+4)}$. Similarly, for the semi-supervised classifier, we set

$$\hat{k}_{\mathrm{SS}}(x) := \max\Big[1, \min\big[\lfloor \hat{B}_{\mathrm{SS}}\{\hat{f}_m(x)n/\|\hat{f}_m\|_\infty\}^{4/(d+4)} \rfloor, n/2\big]\Big],$$

where $\hat{B}_{\mathrm{SS}}$ was chosen analogously to $\hat{B}_{\mathrm{O}}$, and where $\hat{f}_m$ is the d-dimensional kernel density estimator constructed using a truncated normal kernel and bandwidths chosen via the default method in the `R` package `ks` (Duong, 2015). In practice, we estimated $\|\hat{f}_m\|_\infty$ by the maximum value attained on the unlabelled training set.

In each of the three settings above, we generated a training set of size $n \in \{50, 200, 1000\}$ in dimensions $d \in \{1, 2, 5\}$, an unlabelled training set of size 1000, and a test set of size 1000. In Table 1, we present the sample mean and standard error (in subscript) of the risks computed from 1000 repetitions of each experiment. Further, we present estimates of the regret ratios, given by

$$\frac{R(\hat{C}_n^{\hat{k}_{\mathrm{O}}\mathrm{nn}}) - R(C^{\mathrm{Bayes}})}{R(\hat{C}_n^{\hat{k}\mathrm{nn}}) - R(C^{\mathrm{Bayes}})} \quad \text{and} \quad \frac{R(\hat{C}_n^{\hat{k}_{\mathrm{SS}}\mathrm{nn}}) - R(C^{\mathrm{Bayes}})}{R(\hat{C}_n^{\hat{k}\mathrm{nn}}) - R(C^{\mathrm{Bayes}})}, \tag{7}$$

for which the standard errors given are estimated via the delta method. From Table 1, we saw improvement in performance from the oracle and semi-supervised classifiers in 22 of the 27 experiments, comparable performance in three experiments, and there were two where the standard $k$nn classifier was the best of the three classifiers considered. In those latter two cases, the theoretical improvement expected for the local classifiers is small; for instance, when $d = 5$ in Setting 2, the excess risk for the local classifiers converges at rate $O(n^{-4/9})$, while the standard k-nearest neighbour classifier can attain a rate at least as fast as $o(n^{-1/3+\epsilon})$ for every $\epsilon > 0$. It is therefore perhaps unsurprising that we require the larger sample size of $n = 1000$ for the local classifiers to yield an improvement in this case. The semi-supervised classifier exhibits similar performance to the oracle classifier in all settings, though some deterioration is noticeable in higher dimensions, where it is harder to construct a good estimate of $\bar{f}$ from the unlabelled training data.

6 Proofs

In this section, we provide proofs of all of our claimed results, which rely on the general asymptotic expansion presented in Theorem 4 below. We begin with some further notation. Define the $d \times n$ matrices $X^n := (X_1 \dots X_n)$ and $x^n := (x_1 \dots x_n)$. Write

$$\hat{\mu}_n(x) = \hat{\mu}_n(x, x^n) := \mathbb{E}\{\hat{S}_n(x) | X^n = x^n\} = \frac{1}{k_{\mathrm{L}}(x)} \sum_{i=1}^{k_{\mathrm{L}}(x)} \eta(x_{(i)}),$$

and

$$\hat{\sigma}_n^2(x) = \hat{\sigma}_n^2(x, x^n) := \mathrm{Var}\{\hat{S}_n(x) | X^n = x^n\} = \frac{1}{k_{\mathrm{L}}(x)^2} \sum_{i=1}^{k_{\mathrm{L}}(x)} \eta(x_{(i)})\{1 - \eta(x_{(i)})\}.$$

Here we have used the fact that the ordered labels $Y_{(1)}, \dots, Y_{(n)}$ are independent given X^n, satisfying $\mathbb{P}(Y_{(i)} = 1 | X^n) = \eta(X_{(i)})$. Since η takes values in $[0, 1]$ it is clear that $0 \leq \hat{\sigma}_n^2(x) \leq \frac{1}{4k_{\mathrm{L}}(x)}$ for all $x \in \mathbb{R}^d$. Further, write $\mu_n(x) := \mathbb{E}\{\hat{S}_n(x)\} = \frac{1}{k_{\mathrm{L}}(x)} \sum_{i=1}^{k_{\mathrm{L}}(x)} \mathbb{E}\eta(X_{(i)})$ for the unconditional expectation of $\hat{S}_n(x)$. Recall also that $p_r(x) = P_X\big(B_r(x)\big)$.

6.1 A general asymptotic expansion

Let $c_n := \sup_{x_0 \in \mathcal{S} : \bar{f}(x_0) \geq k_{\mathrm{L}}(x_0)/(n-1)} g(x_0)$, where g is defined in assumption **(A.2)**, and for $x \in \mathbb{R}^d$, let

$$\delta_n(x) = \delta_{n,\mathrm{L}}(x) := \frac{k_{\mathrm{L}}(x)}{n-1} c_n^d \log^d\Big(\frac{n-1}{k_{\mathrm{L}}(x)}\Big). \tag{8}$$

Recall that $\mathcal{S} = \{x \in \mathbb{R}^d : \eta(x) = 1/2\}$, and note that by Proposition 5 in the Appendix, for $\epsilon > 0$, we can write

$$\mathcal{S}^\epsilon = \left\{ x_0 + t \frac{\dot{\eta}(x_0)}{\|\dot{\eta}(x_0)\|} : x_0 \in \mathcal{S}, |t| < \epsilon \right\}.$$

Let

$$\epsilon_n := \frac{1}{c_n \beta^{1/2} \log^{1/2}(n-1)}, \tag{9}$$

and recall the definition of the function $a(\cdot)$ in (3).

Theorem 4. *Assume **(A.1)**, **(A.2)**, **(A.3)** and **(A.4)(ρ)**, for some $\rho > 0$. For n sufficiently large, let $\mathcal{R}_n \subseteq \big\{x \in \mathbb{R}^d : \bar{f}(x) \geq \delta_n(x)\big\}$ be a d-dimensional manifold. Write $\partial\mathcal{R}_n$ for the topological boundary of $\mathcal{R}_n$, let $(\partial\mathcal{R}_n)^\epsilon := \partial\mathcal{R}_n + \epsilon \bar{B}_1(0)$, and let $\mathcal{S}_n := \mathcal{S} \cap \mathcal{R}_n$. For $\beta \in (0, 1/2)$ and $\tau > 0$ define the class of functions*

$$K_{\beta,\tau} := \left\{ k_{\mathrm{L}} : \mathbb{R}^d \to K_\beta : \sup_{x_0 \in \mathcal{S}_n} \sup_{|t| < \epsilon_n} \left| \frac{k_{\mathrm{L}}\big(x_0 + t \frac{\dot{\eta}(x_0)}{\|\dot{\eta}(x_0)\|}\big)}{k_{\mathrm{L}}(x_0)} - 1 \right| \leq \tau \right\}.$$

Then for each $\beta \in (0, 1/2)$ and each $\tau = \tau_n$ with $\tau_n \searrow 0$, we have

$$R_{\mathcal{R}_n}(\hat{C}_n^{k_{\mathrm{L}}\mathrm{nn}}) - R_{\mathcal{R}_n}(C^{\mathrm{Bayes}}) = \int_{\mathcal{S}_n} \frac{\bar{f}(x_0)}{\|\dot{\eta}(x_0)\|} \left\{ \frac{1}{4k_{\mathrm{L}}(x_0)} + \Big(\frac{k_{\mathrm{L}}(x_0)}{n\bar{f}(x_0)}\Big)^{4/d} a(x_0)^2 \right\} d\mathrm{Vol}^{d-1}(x_0) + o\big(\gamma_n(k_{\mathrm{L}})\big) + O\big\{P_X\big((\partial\mathcal{R}_n)^{\epsilon_n} \cap \mathcal{S}^{\epsilon_n}\big)\big\}$$

as $n \to \infty$, uniformly for $k_{\mathrm{L}} \in K_{\beta,\tau}$, where

$$\gamma_n(k_{\mathrm{L}}) := \int_{\mathcal{S}_n} \frac{\bar{f}(x_0)}{\|\dot{\eta}(x_0)\|} \left\{ \frac{1}{4k_{\mathrm{L}}(x_0)} + \Big(\frac{k_{\mathrm{L}}(x_0)}{n\bar{f}(x_0)}\Big)^{4/d} g(x_0)^2 \right\} d\mathrm{Vol}^{d-1}(x_0).$$

Proof of Theorem 4. First observe that

$$R_{\mathcal{R}_n}(\hat{C}_n^{k_{\mathrm{L}}\mathrm{nn}}) - R_{\mathcal{R}_n}(C^{\mathrm{Bayes}}) = \int_{\mathcal{R}_n} \big[\mathbb{P}\{\hat{S}_n(x) < 1/2\} - \mathbb{1}_{\{\eta(x)<1/2\}}\big]\{2\eta(x) - 1\}\bar{f}(x)\, dx. \quad (10)$$

The proof is presented in seven steps. We will see that the dominant contribution to the integral in (10) arises from a small neighbourhood about the Bayes decision boundary, i.e. the region $\mathcal{S}^{\epsilon_n} \cap \mathcal{R}_n$. On $\mathcal{R}_n \setminus \mathcal{S}^{\epsilon_n}$, the k_{L}nn classifier agrees with the Bayes classifier with high probability (asymptotically). More precisely, we show in Step 4 that

$$\sup_{k_{\mathrm{L}} \in K_{\beta,\tau}} \sup_{x \in \mathcal{R}_n \setminus \mathcal{S}^{\epsilon_n}} |\mathbb{P}\{\hat{S}_n(x) < 1/2\} - \mathbb{1}_{\{\eta(x)<1/2\}}| = O(n^{-M}),$$

for each $M > 0$, as $n \to \infty$. In Steps 1, 2 and 3, we derive the key asymptotic properties of the bias, conditional (on X^n) bias and variance of $\hat{S}_n(x)$ respectively. In Step 5 we show that the integral over $\mathcal{S}^{\epsilon_n} \cap \mathcal{R}_n$ can be decomposed into an integral over $\mathcal{S}_n$ and one perpendicular to $\mathcal{S}$. Step 6 is dedicated to combining the results of Steps 1 - 5; we derive the leading order terms in the asymptotic expansion of the integral in (10). Finally, we bound the remaining error terms to conclude the proof in Step 7. To ease notation, where it is clear from the context, we write k_{L} in place of $k_{\mathrm{L}}(x)$.

Step 1: Let $\mu_n(x) := \mathbb{E}\{\hat{S}_n(x)\}$, and for $x_0 \in \mathcal{S}$ and $t \in \mathbb{R}$, write $x = x(x_0, t) := x_0 + t\frac{\dot{\eta}(x_0)}{\|\dot{\eta}(x_0)\|}$. We show that

$$\mu_n(x) - \eta(x) - \Big(\frac{k_{\mathrm{L}}(x)}{n\bar{f}(x)}\Big)^{2/d} a(x) = o\bigg(\Big(\frac{k_{\mathrm{L}}(x_0)}{n\bar{f}(x_0)}\Big)^{2/d} g(x_0)\bigg),$$

uniformly for $k_{\mathrm{L}} \in K_{\beta,\tau}$, $x_0 \in \mathcal{S}_n$ and $|t| < \epsilon_n$. Write

$$\begin{aligned}\mu_n(x) &= \eta(x) + \frac{1}{k_{\mathrm{L}}(x)} \sum_{i=1}^{k_{\mathrm{L}}(x)} \mathbb{E}\{\eta(X_{(i)}) - \eta(x)\} \\ &= \eta(x) + \frac{1}{k_{\mathrm{L}}(x)} \sum_{i=1}^{k_{\mathrm{L}}(x)} \mathbb{E}\{(X_{(i)} - x)^T\dot{\eta}(x)\} + \frac{1}{2}\mathbb{E}\{(X_{(i)} - x)^T\ddot{\eta}(x)(X_{(i)} - x)\} + R_1,\end{aligned}$$

where we show in Step 7 that

$$|R_1| = o\left\{\left(\frac{k_{\mathrm{L}}(x_0)}{n\bar{f}(x_0)}\right)^{2/d}\right\} \tag{11}$$

uniformly for $k_{\mathrm{L}} \in K_{\beta,\tau}$, $x_0 \in \mathcal{S}_n$ and $|t| < \epsilon_n$.

The density of $X_{(i)} - x$ at $u \in \mathbb{R}^d$ is given by

$$f_{(i)}(u) := n\bar{f}(x+u)\binom{n-1}{i-1}p_{\|u\|}^{i-1}(1-p_{\|u\|})^{n-i} = n\bar{f}(x+u)p_{\|u\|}^{n-1}(i-1), \tag{12}$$

where $p_{\|u\|} = p_{\|u\|}(x)$ and $p_{\|u\|}^{n-1}(i-1)$ denotes the probability that a $\mathrm{Bin}(n-1, p_{\|u\|})$ random variable equals $i-1$. Now let

$$r_n = r_n(x) := \left\{\frac{2k_{\mathrm{L}}(x)}{(n-1)\bar{f}(x)a_d}\right\}^{1/d}. \tag{13}$$

We show in Step 7 that

$$R_2 := \sup_{k_{\mathrm{L}} \in K_{\beta,\tau}} \sup_{x_0 \in \mathcal{S}_n} \sup_{|t|<\epsilon_n} \mathbb{E}\{\|X_{(k_{\mathrm{L}})} - x\|^2 \mathbb{1}_{\{\|X_{(k_{\mathrm{L}})}-x\| \geq r_n\}}\} = O(n^{-M}), \tag{14}$$

for each $M > 0$, as $n \to \infty$. It follows from (12) and (14), together with the assumption on $\|\dot{\eta}(\cdot)\|$ in **(A.3)** that

$$\mathbb{E}\{(X_{(i)} - x)^T\dot{\eta}(x)\} = \int_{B_{r_n}(0)} \dot{\eta}(x)^T un\{\bar{f}(x+u) - \bar{f}(x)\}p_{\|u\|}^{n-1}(i-1)\,du + O(n^{-M}),$$

uniformly for $1 \leq i \leq k_{\mathrm{L}}$, $x_0 \in \mathcal{S}_n$ and $|t| < \epsilon_n$. Similarly, using the assumption on $\|\ddot{\eta}(\cdot)\|_{\mathrm{op}}$ in **(A.3)**,

$$\mathbb{E}\{(X_{(i)} - x)^T\ddot{\eta}(x)(X_{(i)} - x)\} = \int_{B_{r_n}(0)} u^T\ddot{\eta}(x)un\bar{f}(x+u)p_{\|u\|}^{n-1}(i-1)\,du + O(n^{-M}),$$

uniformly for $1 \leq i \leq k_{\mathrm{L}}$, $x_0 \in \mathcal{S}_n$ and $|t| < \epsilon_n$. Hence, summing over i, we see that

$$\begin{aligned}
&\frac{1}{k_{\mathrm{L}}}\sum_{i=1}^{k_{\mathrm{L}}}\mathbb{E}\{(X_{(i)} - x)^T\dot{\eta}(x)\} + \frac{1}{2k_{\mathrm{L}}}\sum_{i=1}^{k_{\mathrm{L}}}\mathbb{E}\{(X_{(i)} - x)^T\ddot{\eta}(x)(X_{(i)} - x)\} \\
&= \int_{B_{r_n}(0)}\left[\dot{\eta}(x)^T un\{\bar{f}(x+u) - \bar{f}(x)\} + \frac{1}{2}u^T\ddot{\eta}(x)un\bar{f}(x+u)\right]q_{\|u\|}^{n-1}(k_{\mathrm{L}})\,du + O(n^{-M}),
\end{aligned}$$

where $q_{\|u\|}^{n-1}(k_{\mathrm{L}})$ denotes the probability that a $\mathrm{Bin}(n-1, p_{\|u\|})$ random variable is less than k_{L}. Let $n_0 \in \mathbb{N}$ be large enough that $\epsilon_n + \sup_{x_0 \in \mathcal{S}_n}\sup_{|t|<\epsilon_n} r_n(x) < \epsilon_0$ for $n \geq n_0$. That this is possible follows from the fact that, for $\epsilon_n < \epsilon_0$,

$$\sup_{k_{\mathrm{L}} \in K_{\beta,\tau}} \sup_{x_0 \in \mathcal{S}_n} \sup_{|t|<\epsilon_n} \max\left\{\left|\frac{k_{\mathrm{L}}(x)}{k_{\mathrm{L}}(x_0)} - 1\right|, \left|\frac{\bar{f}(x)}{\bar{f}(x_0)} - 1\right|\right\} \leq \max\{\tau, c_n\epsilon_n + c_n\epsilon_n^2/2\} \to 0. \tag{15}$$

By a Taylor expansion of $\bar{f}$ and assumption **(A.2)**, for all $x_0 \in \mathcal{S}_n$, $|t| < \epsilon_n$, $\|u\| < r_n$ and $n \geq n_0$,

$$\left|\bar{f}(x+u) - \bar{f}(x) - u^T\dot{\bar{f}}(x)\right| \leq \frac{\|u\|^2}{2}\sup_{s\in B_{\|u\|}(0)}\|\ddot{\bar{f}}(x+s)\|_{\mathrm{op}} \leq \frac{\|u\|^2}{2}\bar{f}(x_0)g(x_0).$$

Hence, for $x_0 \in \mathcal{S}_n$, $|t| < \epsilon_n$, $r < r_n$ and $n \geq n_0$,

$$\begin{aligned}|p_r(x) - \bar{f}(x)a_d r^d| &\leq \int_{B_r(0)} |\bar{f}(x+u) - \bar{f}(x) - u^T\dot{\bar{f}}(x)|\,du \\ &\leq \frac{1}{2}\bar{f}(x_0)g(x_0)\int_{B_r(0)}\|u\|^2\,du = \frac{da_d}{2(d+2)}\bar{f}(x_0)g(x_0)r^{d+2}. \end{aligned} \tag{16}$$

Now, for $v \in B_1(0)$, $x_0 \in \mathcal{S}_n$, $|t| < \epsilon_n$ and $n \geq n_0$,

$$k_{\mathrm{L}}(x) - (n-1)p_{\|v\|r_n} = k_{\mathrm{L}}(x) - (n-1)\bar{f}(x)a_d\|v\|^d r_n^d + R_3 = k_{\mathrm{L}}(x)(1-2\|v\|^d) + R_3,$$

where

$$\begin{aligned}|R_3| &\leq \frac{da_d(n-1)\bar{f}(x_0)g(x_0)\|v\|^{d+2}r_n^{d+2}}{2(d+2)} \\ &= \frac{dk_{\mathrm{L}}(x)\bar{f}(x_0)g(x_0)\|v\|^{d+2}r_n^2}{(d+2)\bar{f}(x)} \leq \frac{2^{2/d}dk_{\mathrm{L}}(x)}{a_d^{2/d}(d+2)\log^2\left(\frac{n-1}{k_{\mathrm{L}}(x_0)}\right)}\left(\frac{\bar{f}(x_0)}{\bar{f}(x)}\right)^{1+2/d}\left(\frac{k_{\mathrm{L}}(x)}{k_{\mathrm{L}}(x_0)}\right)^{2/d}.\end{aligned}$$

It follows from (15) that there exists $n_1 \in \mathbb{N}$ such that, for all $x_0 \in \mathcal{S}_n$, $|t| < \epsilon_n$, $\|v\|^d \in (0, 1/2 - 1/\log\left((n-1)/k_{\mathrm{L}}(x_0)\right)]$ and $n \geq n_1$,

$$k_{\mathrm{L}}(x) - (n-1)p_{\|v\|r_n} \geq \frac{k_{\mathrm{L}}(x)}{\log((n-1)/k_{\mathrm{L}}(x_0))},$$

Similarly, for all $\|v\|^d \in [1/2 + 1/\log((n-1)/k_{\mathrm{L}}(x_0)), 1)$ and $n \geq n_1$,

$$(n-1)p_{\|v\|r_n} - k_{\mathrm{L}}(x) \geq \frac{k_{\mathrm{L}}(x)}{\log((n-1)/k_{\mathrm{L}}(x_0))}.$$

Hence, by Bernstein's inequality, we have that for each $M > 0$,

$$\sup_{k_{\mathrm{L}}\in K_{\beta,\tau}}\sup_{x_0\in\mathcal{S}_n}\sup_{|t|<\epsilon_n}\sup_{\|v\|^d\in(0,1/2-1/\log((n-1)/k_{\mathrm{L}}(x_0))]} 1 - q_{\|v\|r_n}^{n-1}(k_{\mathrm{L}}(x)) = O(n^{-M}),$$

and

$$\sup_{k_{\mathrm{L}}\in K_{\beta,\tau}}\sup_{x_0\in\mathcal{S}_n}\sup_{|t|<\epsilon_n}\sup_{\|v\|^d\in[1/2+1/\log((n-1)/k_{\mathrm{L}}(x_0)),1)} q_{\|v\|r_n}^{n-1}(k_{\mathrm{L}}(x)) = O(n^{-M}). \tag{17}$$

We conclude that

$$\begin{aligned}
&\frac{1}{k_{\mathrm{L}}(x)} \int_{B_{r_n}(0)} \Big[\dot{\eta}(x)^T u n \{ \bar{f}(x+u) - \bar{f}(x) \} + \frac{1}{2} u^T \ddot{\eta}(x) u n \bar{f}(x+u) \Big] q_{\|u\|}^{n-1}(k_{\mathrm{L}}(x))\, du \\
&\quad = \frac{1}{k_{\mathrm{L}}(x)} \int_{B_{2^{-1/d} r_n}(0)} \Big[\dot{\eta}(x)^T u n \{ \bar{f}(x+u) - \bar{f}(x) \} + \frac{1}{2} u^T \ddot{\eta}(x) u n \bar{f}(x+u) \Big]\, du + R_{41} \\
&\quad = \Big(\frac{k_{\mathrm{L}}(x)}{n} \Big)^{2/d} \frac{\sum_{j=1}^d \{ \eta_j(x) \bar{f}_j(x) + \frac{1}{2} \eta_{jj}(x) \bar{f}(x) \}}{(d+2) a_d^{2/d} \bar{f}(x)^{1+2/d}} + R_{41} + R_{42} \\
&\quad = \Big(\frac{k_{\mathrm{L}}(x)}{n \bar{f}(x)} \Big)^{2/d} a(x) + R_{41} + R_{42}, \qquad (18)
\end{aligned}$$

where

$$|R_{41}| + |R_{42}| = o\bigg(\Big(\frac{k_{\mathrm{L}}(x_0)}{n \bar{f}(x_0)} \Big)^{2/d} g(x_0) \bigg),$$

uniformly for $k_{\mathrm{L}} \in K_{\beta,\tau}$, $x_0 \in \mathcal{S}_n$ and $|t| < \epsilon_n$.

Step 2: Recall that $\hat{\sigma}_n^2(x, x^n) = \mathrm{Var}\{ \hat{S}_n(x) | X^n = x^n \}$. We show that

$$\Big| \hat{\sigma}_n^2(x, X^n) - \frac{1}{4 k_{\mathrm{L}}} \Big| = o_p(1/k_{\mathrm{L}}), \qquad (19)$$

uniformly for $k_{\mathrm{L}} \in K_{\beta,\tau}$, $x_0 \in \mathcal{S}_n$ and $|t| < \epsilon_n$. Recall that

$$\hat{\sigma}_n^2(x, X^n) = \frac{1}{k_{\mathrm{L}}^2} \sum_{i=1}^{k_{\mathrm{L}}} \eta(X_{(i)}) \{ 1 - \eta(X_{(i)}) \}.$$

Let $n_2 \in \mathbb{N}$ be large enough that $1 - c_n \epsilon_n - \frac{d+1}{d+2} c_n \epsilon_n^2 \geq \mu_0 / a_d$ for $n \geq n_2$. Then for $n \geq \max\{n_0, n_2\}$, $\epsilon < \epsilon_n$, $x_0 \in \mathcal{S}_n$ and $|t| < \epsilon_n$, we have by **(A.2)** and a very similar argument to that in (16) that

$$p_\epsilon(x) \geq \mu_0 \epsilon^d \bar{f}(x_0) \geq \mu_0 \epsilon^d \delta_n(x_0). \qquad (20)$$

Now suppose that $z_1, \ldots, z_N \in \mathcal{R}_n \cup \mathcal{S}_n^{\epsilon_n}$ are such that $\|z_j - z_\ell\| \geq \epsilon_n / 6$ for all $j \neq \ell$, but $\sup_{x \in \mathcal{R}_n \cup \mathcal{S}_n^{\epsilon_n}} \min_{j=1,\ldots,N} \|x - z_j\| < \epsilon_n / 6$. We have by **(A.2)** that

$$1 = P_X(\mathbb{R}^d) \geq \sum_{j=1}^N p_{\epsilon_n/12}(z_j) \geq \frac{N \mu_0 \beta^{d/2} \log^{d/2}(n-1)}{12^d (n-1)^{1-\beta}}.$$

For each $j = 1, \ldots, N$, choose

$$z_j' \in \operatorname*{argmax}_{z \in B_{z_j}(\epsilon_n/6) \cap (\mathcal{R}_n \cup \mathcal{S}_n^{\epsilon_n})} k_{\mathrm{L}}(z).$$

Now, given $x \in \mathcal{R}_n \cup \mathcal{S}_n^{\epsilon_n}$, let $j_0 := \mathrm{argmin}_j \|x - z_j\|$, so that $B_{\epsilon_n/6}(z_{j_0}') \subseteq B_{\epsilon_n/2}(x)$. Thus, if there are at least $k_{\mathrm{L}}(z_j')$ points among $\{x_1, \ldots, x_n\}$ inside each of the balls $B_{\epsilon_n/6}(z_j')$, then

for every $x \in \mathcal{R}_n \cup \mathcal{S}_n^{\epsilon_n}$ there are at least $k_{\mathrm{L}}(x)$ of them in $B_{\epsilon_n/2}(x)$. Moreover by (15), (20) and **(A.2)**

$$\min_{j=1,\dots,N}\Big\{np_{\epsilon_n/6}(z_j') - 2k_{\mathrm{L}}(z_j')\Big\} \geq (n-1)^{\beta}$$

for all $k_{\mathrm{L}} \in K_{\beta,\tau}$ and $n \geq n_3$, say. Define $A_{k_{\mathrm{L}}} := \big\{\|X_{(k_{\mathrm{L}})}(x) - x\| < \epsilon_n/2 \text{ for all } x \in \mathcal{R}_n \cup \mathcal{S}_n^{\epsilon_n}\big\}$. Then by a standard binomial tail bound (Shorack and Wellner, 1986, Equation (6), p. 440), for $n \geq n_3$ and any $M > 0$,

$$\begin{aligned}
\mathbb{P}(A_{k_{\mathrm{L}}}^c) &= \mathbb{P}\Big\{\sup_{x \in \mathcal{R}_n \cup \mathcal{S}_n^{\epsilon_n}} \|X_{(k_{\mathrm{L}}(x))}(x) - x\| \geq \epsilon_n/2\Big\} \\
&\leq \mathbb{P}\Big\{\max_{j=1,\dots,N} \|X_{(k_{\mathrm{L}}(z_j))}(z_j') - z_j'\| \geq \epsilon_n/6\Big\} \leq \sum_{j=1}^{N} \mathbb{P}\big\{\|X_{(k_{\mathrm{L}}(z_j))}(z_j') - z_j'\| \geq \epsilon_n/6\big\} \\
&\leq N \max_{j=1,\dots,N} \exp\Big(-\frac{1}{2}np_{\epsilon_n/6}(z_j') + k_{\mathrm{L}}(z_j')\Big) = O(n^{-M}),
\end{aligned} \tag{21}$$

uniformly for $k_{\mathrm{L}} \in K_{\beta,\tau}$. Now,

$$\sup_{k_{\mathrm{L}} \in K_{\beta,\tau}} \sup_{x_0 \in \mathcal{S}_n} \sup_{|t| < \epsilon_n} \sup_{x^n \in A_{k_{\mathrm{L}}}} \max_{1 \leq i \leq k_{\mathrm{L}}(x)} |\eta(x_{(i)}(x)) - 1/2| \to 0.$$

It follows that

$$\sup_{x^n \in A_{k_{\mathrm{L}}}} \Big|\frac{1}{k_{\mathrm{L}}(x)^2} \sum_{i=1}^{k_{\mathrm{L}}(x)} \eta(x_{(i)}(x))\{1 - \eta(x_{(i)}(x))\} - \frac{1}{4k_{\mathrm{L}}(x)}\Big| = o\Big(\frac{1}{k_{\mathrm{L}}(x)}\Big) \tag{22}$$

as $n \to \infty$, uniformly for $x_0 \in \mathcal{S}_n$, $|t| < \epsilon_n$ and $k_{\mathrm{L}} \in K_{\beta,\tau}$. The claim (19) follows from (21) and (22).

Step 3: In this step, we emphasise the dependence of $\hat{\mu}_n(x, x^n) = \mathbb{E}\{\hat{S}_n(x) | X^n = x^n\}$ on k_{L} by writing it as $\hat{\mu}_n^{(k_{\mathrm{L}})}(x, x^n)$. We show that

$$\mathrm{Var}\{\hat{\mu}_n^{(k_{\mathrm{L}})}(x, X^n)\} = O\Big\{\frac{1}{k_{\mathrm{L}}}\Big(\frac{k_{\mathrm{L}}(x_0)}{n\bar{f}(x_0)}\Big)^{2/d}\Big\} \tag{23}$$

uniformly for $x_0 \in \mathcal{S}_n$, $|t| < \epsilon_n$ and $k_{\mathrm{L}} \in K_{\beta,\tau}$. We write $X^{n,j} := (X_1 \dots X_{j-1}\ X_{j+1} \dots X_n)$, considered as a random $d \times (n-1)$ matrix, so that

$$\hat{\mu}_n^{(k_{\mathrm{L}})}(x, X^n) - \hat{\mu}_{n-1}^{(k_{\mathrm{L}})}(x, X^{n,(i)}) = \frac{1}{k_{\mathrm{L}}}\{\eta(X_{(i)}) - \eta(X_{(k_{\mathrm{L}}+1)})\}\mathbb{1}_{\{i \leq k_{\mathrm{L}}\}}.$$

It follows from the Efron–Stein inequality (e.g. Boucheron, Lugosi and Massart, 2013, The-

orem 3.1) that

$$\begin{aligned}\mathrm{Var}\{\hat{\mu}_n^{(k_{\mathrm{L}})}(x, X^n)\} &\leq \sum_{i=1}^{n} \mathbb{E}\big[\{\hat{\mu}_n^{(k_{\mathrm{L}})}(x, X^n) - \hat{\mu}_{n-1}^{(k_{\mathrm{L}})}(x, X^{n,(i)})\}^2\big] \\ &= \frac{1}{k_{\mathrm{L}}^2} \sum_{i=1}^{k_{\mathrm{L}}} \mathbb{E}\big[\{\eta(X_{(i)}) - \eta(X_{(k_{\mathrm{L}}+1)})\}^2\big] \\ &\leq \frac{2}{k_{\mathrm{L}}^2} \sum_{i=1}^{k_{\mathrm{L}}} \mathbb{E}\big[\{\eta(X_{(i)}) - \eta(x)\}^2 + \{\eta(X_{(k_{\mathrm{L}}+1)}) - \eta(x)\}^2\big]. \qquad (24)\end{aligned}$$

Recall the definition of r_n given in (13). Now observe that, for $\max(\epsilon_n, r_n) \leq \epsilon_0$ and all $M > 0$ we have that

$$\begin{aligned}&\max_{i \in \{1,\ldots,k_{\mathrm{L}}+1\}} \mathbb{E}\big[\{\eta(X_{(i)}) - \eta(x)\}^2\big] \\ &\quad\leq \max_{i \in \{1,\ldots,k_{\mathrm{L}}+1\}} \mathbb{E}\big[\{\eta(X_{(i)}) - \eta(x)\}^2 \mathbb{1}_{\{\|X_{(i)} - x\| \leq r_n\}}\big] + \mathbb{P}(\|X_{(k_{\mathrm{L}}+1)} - x\| > r_n) \\ &\quad\leq r_n^2 \sup_{z \in \mathcal{S}^{2\epsilon_0}} \|\dot{\eta}(z)\|^2 + O(n^{-M}), \qquad (25)\end{aligned}$$

uniformly for $x_0 \in \mathcal{S}_n$, $|t| < \epsilon_n$ and $k_{\mathrm{L}} \in K_{\beta,\tau}$. The final inequality here follows from similar arguments to those used to bound R_1. Now (23) follows from (24) and (25).

Step 4: We show that

$$\sup_{k_{\mathrm{L}} \in K_{\beta,\tau}} \sup_{x \in \mathcal{R}_n \setminus \mathcal{S}^{\epsilon_n}} |\mathbb{P}\{\hat{S}_n(x) < 1/2\} - \mathbb{1}_{\{\eta(x) < 1/2\}}| = O(n^{-M}),$$

for each $M > 0$, as $n \to \infty$. First, by **(A.3)** and Proposition 5 in Section 7.2, there exists $c_0 > 0$ such that for every $\epsilon \in (0, \epsilon_0]$,

$$\inf_{x \in \mathcal{R}_n \setminus \mathcal{S}^{\epsilon}} |\eta(x) - 1/2| \geq c_0 \min\Big\{\epsilon\,, \inf_{x \in \mathcal{R}_n \setminus \mathcal{S}^{\epsilon_0}} \delta_n(x)^{\beta/2}\Big\}.$$

Hence, on the event $A_{k_{\mathrm{L}}}$, for $\epsilon_n < \epsilon_0$ and $x \in \mathcal{R}_n \setminus \mathcal{S}^{\epsilon_n}$, all of the k_{L} nearest neighbours of x are on the same side of $\mathcal{S}$, so

$$\begin{aligned}|\hat{\mu}_n(x, X^n) - 1/2| = \Big|\frac{1}{k_{\mathrm{L}}} \sum_{i=1}^{k_{\mathrm{L}}} \eta(X_{(i)}) - 1/2\Big| &\geq \inf_{z \in B_{\epsilon_n/2}(x)} |\eta(z) - 1/2| \\ &\geq c_0 \min\Big\{\frac{\epsilon_n}{2}\,, \inf_{x \in \mathcal{R}_n \setminus \mathcal{S}^{\epsilon_0}} \delta_n(x)^{\beta/2}\Big\}.\end{aligned}$$

Now, conditional on X^n, $\hat{S}_n(x)$ is the sum of $k_{\mathrm{L}}(x)$ independent terms. Therefore, by Hoeffding's inequality,

$$\begin{aligned}\sup_{x\in\mathcal{R}_n\setminus\mathcal{S}^{\epsilon_n}}\bigl|\mathbb{P}\{\hat{S}_n(x)<1/2\}-\mathbb{1}_{\{\eta(x)\le 1/2\}}\bigr| &= \sup_{x\in\mathcal{R}_n\setminus\mathcal{S}^{\epsilon_n}}\bigl|\mathbb{E}\{\mathbb{P}\{\hat{S}_n(x)<1/2|X^n\}-\mathbb{1}_{\{\eta(x)\le 1/2\}}\bigr| \\ &\le \sup_{x\in\mathcal{R}_n\setminus\mathcal{S}^{\epsilon_n}}\mathbb{E}\bigl[\exp(-2k_{\mathrm{L}}\{\hat{\mu}_n(x,X^n)-1/2\}^2)\mathbb{1}_{A_{k_{\mathrm{L}}}}\bigr]+\mathbb{P}(A_{k_{\mathrm{L}}}^c)=O(n^{-M})\end{aligned}$$

for every $M>0$. This completes Step 4.

Step 5: It is now convenient to be more explicit in our notation, by writing $x_0^t := x_0+t\dot{\eta}(x_0)/\|\dot{\eta}(x_0)\|$. We also let

$$\psi(x) := \{2\eta(x)-1\}\bar{f}(x) = \pi_1 f_1(x)-\pi_2 f_2(x).$$

Recalling that $\mathcal{S}_n := \mathcal{S}\cap\mathcal{R}_n$, we show that

$$\begin{aligned}&\int_{\mathcal{S}^{\epsilon_n}\cap\mathcal{R}_n}\psi(x)[\mathbb{P}\{\hat{S}_n(x)<1/2\}-\mathbb{1}_{\{\eta(x)<1/2\}}]\,dx \\ &=\int_{\mathcal{S}_n}\int_{-\epsilon_n}^{\epsilon_n}\psi(x_0^t)[\mathbb{P}\{\hat{S}_n(x_0^t)<1/2\}-\mathbb{1}_{\{t<0\}}]\,dt\,d\mathrm{Vol}^{d-1}(x_0)\{1+o(1)\} \\ &\qquad\qquad + O\bigl\{P_X\bigl((\partial\mathcal{R}_n)^{\epsilon_n}\cap\mathcal{S}^{\epsilon_n}\bigr)\bigr\},\end{aligned}$$

uniformly for $k_{\mathrm{L}}\in K_{\beta,\tau}$. Now by Proposition 6 in Section 7.2, for

$$\epsilon_n \le \min\left\{\epsilon_0\,,\ \frac{\inf_{x_0\in\mathcal{S}}\|\dot{\eta}(x_0)\|}{\sup_{z\in\mathcal{S}^{\epsilon_0}}\|\ddot{\eta}(z)\|_{\mathrm{op}}}\right\},$$

the map $x(x_0,t) := x_0^t$ is a diffeomorphism from $\mathcal{S}_n\times(-\epsilon_n,\epsilon_n)$ to $\mathcal{S}_n^{\epsilon_n}$, where

$$\mathcal{S}_n^{\epsilon} := \left\{x_0+t\frac{\dot{\eta}(x_0)}{\|\dot{\eta}(x_0)\|} : x_0\in\mathcal{S}_n, |t|<\epsilon\right\}.$$

Furthermore, for such n, and $|t|<\epsilon_n$, $\mathrm{sgn}\{\eta(x_0^t)-1/2\}=\mathrm{sgn}(t)$. Now observe that $(\mathcal{S}^{\epsilon_n}\cap\mathcal{R}_n)\setminus\mathcal{S}_n^{\epsilon_n}\subseteq(\partial\mathcal{R}_n)^{\epsilon_n}\cap\mathcal{S}^{\epsilon_n}$ and $\mathcal{S}_n^{\epsilon_n}\setminus(\mathcal{S}^{\epsilon_n}\cap\mathcal{R}_n)\subseteq(\partial\mathcal{R}_n)^{\epsilon_n}\cap\mathcal{S}^{\epsilon_n}$. It follows from this and (56) in Section 7.3 that

$$\begin{aligned}&\int_{\mathcal{S}^{\epsilon_n}\cap\mathcal{R}_n}\psi(x)[\mathbb{P}\{\hat{S}_n(x)<1/2\}-\mathbb{1}_{\{\eta(x)<1/2\}}]\,dx \\ &=\int_{\mathcal{S}_n^{\epsilon_n}}\psi(x)[\mathbb{P}\{\hat{S}_n(x)<1/2\}-\mathbb{1}_{\{\eta(x)<1/2\}}]\,dx+O\bigl\{P_X\bigl((\partial\mathcal{R}_n)^{\epsilon_n}\cap\mathcal{S}^{\epsilon_n}\bigr)\bigr\} \\ &=\int_{\mathcal{S}_n}\int_{-\epsilon_n}^{\epsilon_n}\det(I+tB)\psi(x_0^t)[\mathbb{P}\{\hat{S}_n(x_0^t)<1/2\}-\mathbb{1}_{\{t<0\}}]\,dt\,d\mathrm{Vol}^{d-1}(x_0) \\ &\qquad\qquad + O\bigl\{P_X\bigl((\partial\mathcal{R}_n)^{\epsilon_n}\cap\mathcal{S}^{\epsilon_n}\bigr)\bigr\}\end{aligned}$$

where B is defined in (49), and $\det(I+tB) = 1+o(1)$ as $n \to \infty$, uniformly for $x_0 \in \mathcal{S}$ and $t \in (-\epsilon_n, \epsilon_n)$.

Step 6: The last step in the main argument is to show that

$$\int_{\mathcal{S}_n} \int_{-\epsilon_n}^{\epsilon_n} \psi(x_0^t)[\mathbb{P}\{\hat{S}_n(x_0^t) < 1/2\} - \mathbb{1}_{\{t<0\}}]\, dt\, d\mathrm{Vol}^{d-1}(x_0)$$
$$= \int_{\mathcal{S}_n} \frac{\bar{f}(x_0)}{\|\dot{\eta}(x_0)\|}\left\{\frac{1}{4k_{\mathrm{L}}(x_0)} + \Big(\frac{k_{\mathrm{L}}(x_0)}{n\bar{f}(x_0)}\Big)^{4/d} a(x_0)^2\right\} d\mathrm{Vol}^{d-1}(x_0) + o(\gamma_n(k_{\mathrm{L}}))$$

as $n \to \infty$, uniformly for $k_{\mathrm{L}} \in K_{\beta,\tau}$. First observe that

$$\int_{\mathcal{S}_n} \int_{-\epsilon_n}^{\epsilon_n} \psi(x_0^t)[\mathbb{P}\{\hat{S}_n(x_0^t) < 1/2\} - \mathbb{1}_{\{t<0\}}]\, dt\, d\mathrm{Vol}^{d-1}(x_0)$$
$$= \int_{\mathcal{S}_n} \int_{-\epsilon_n}^{\epsilon_n} t\|\dot{\psi}(x_0)\|[\mathbb{P}\{\hat{S}_n(x_0^t) < 1/2\} - \mathbb{1}_{\{t<0\}}]\, dt\, d\mathrm{Vol}^{d-1}(x_0)\{1+o(1)\}.$$

Now, write $\mathbb{P}\{\hat{S}_n(x_0^t) < 1/2\} - \mathbb{1}_{\{t<0\}} = \mathbb{E}[\mathbb{P}\{\hat{S}_n(x_0^t) < 1/2 | X^n\} - \mathbb{1}_{\{t<0\}}]$. Note that, given X^n, $\hat{S}_n(x) = \frac{1}{k_{\mathrm{L}}(x)} \sum_{i=1}^{k_{\mathrm{L}}(x)} \mathbb{1}_{\{Y_{(i)}=1\}}$ is the sum of $k_{\mathrm{L}}(x)$ independent Bernoulli variables, satisfying $\mathbb{P}(Y_{(i)} = 1 | X^n) = \eta(X_{(i)})$. Let Φ be the standard normal distribution function, and let

$$\hat{\theta}(x) := -\{\hat{\mu}_n(x, X^n) - 1/2\}/\hat{\sigma}_n(x, X^n)$$
$$\bar{\theta}(x_0, t) := -2k_{\mathrm{L}}(x_0)^{1/2}\left\{t\|\dot{\eta}(x_0)\| + \left(\frac{k_{\mathrm{L}}(x_0)}{n\bar{f}(x_0)}\right)^{2/d} a(x_0)\right\}.$$

We can write

$$\int_{-\epsilon_n}^{\epsilon_n} t\|\dot{\psi}(x_0)\|[\mathbb{P}\{\hat{S}_n(x_0^t) < 1/2\} - \mathbb{1}_{\{t<0\}}]\, dt$$
$$= \int_{-\epsilon_n}^{\epsilon_n} t\|\dot{\psi}(x_0)\|\mathbb{E}\big\{\Phi\big(\hat{\theta}(x_0^t)\big) - \mathbb{1}_{\{t<0\}}\big\}\, dt + R_5(x_0)$$
$$= \int_{-\epsilon_n}^{\epsilon_n} t\|\dot{\psi}(x_0)\|\big\{\Phi\big(\bar{\theta}(x_0, t)\big) - \mathbb{1}_{\{t<0\}}\big\}\, dt + R_5(x_0) + R_6(x_0),$$

where we show in Step 7 that

$$\left|\int_{\mathcal{S}_n} R_5(x_0) + R_6(x_0)\, d\mathrm{Vol}^{d-1}(x_0)\right| = o(\gamma_n(k_{\mathrm{L}})). \tag{26}$$

Then, substituting $u = 2k_{\mathrm{L}}(x_0)^{1/2}t$, we see that

$$\int_{-\epsilon_n}^{\epsilon_n} t\|\dot{\psi}(x_0)\|\big[\Phi\big(\bar{\theta}(x_0, t)\big) - \mathbb{1}_{\{t<0\}}\big]\, dt$$
$$= \frac{1}{4k_{\mathrm{L}}(x_0)} \int_{-2k_{\mathrm{L}}(x_0)^{1/2}\epsilon_n}^{2k_{\mathrm{L}}(x_0)^{1/2}\epsilon_n} u\|\dot{\psi}(x_0)\|\left\{\Phi\left(\bar{\theta}\Big(x_0, \frac{u}{2k_{\mathrm{L}}(x_0)^{1/2}}\Big)\right) - \mathbb{1}_{\{u<0\}}\right\} du$$
$$= \left\{\frac{\bar{f}(x_0)}{4k_{\mathrm{L}}(x_0)\|\dot{\eta}(x_0)\|} + \Big(\frac{k_{\mathrm{L}}(x_0)}{n\bar{f}(x_0)}\Big)^{4/d} \frac{\bar{f}(x_0)a(x_0)^2}{\|\dot{\eta}(x_0)\|}\right\}\{1+o(1)\}.$$

The conclusion follows by integrating with respect to $d\mathrm{Vol}^{d-1}$ over $\mathcal{S}_n$.

Step 7: It remains to bound the error terms R_1, R_2, R_5 and R_6.

To bound R_1: We have

$$R_1 = \frac{1}{k_{\mathrm{L}}} \sum_{i=1}^{k_{\mathrm{L}}} \Big(\mathbb{E}\eta(X_{(i)}) - \eta(x) - \mathbb{E}\{(X_{(i)} - x)^T \dot{\eta}(x)\} - \frac{1}{2}\mathbb{E}\{(X_{(i)} - x)^T \ddot{\eta}(x)(X_{(i)} - x)\} \Big).$$

By a Taylor expansion and the uniform continuity of $\ddot{\eta}$ from **(A.3)**, for all $\epsilon > 0$, there exists $r = r_\epsilon \in (0, \epsilon_0]$, such that for all $x \in \mathcal{S}^{\epsilon_0}$ and $\|z - x\| < r$,

$$\Big| \eta(z) - \eta(x) - (z-x)^T \dot{\eta}(x) - \frac{1}{2}(z-x)^T \ddot{\eta}(x)(z-x) \Big| \leq \epsilon \|z - x\|^2.$$

Hence

$$\begin{aligned} |R_1| \leq \epsilon \frac{1}{k_{\mathrm{L}}} \sum_{i=1}^{k_{\mathrm{L}}} \mathbb{E}\{\|X_{(i)} - x\|^2 \mathbb{1}_{\{\|X_{(k_{\mathrm{L}})} - x\| \leq r\}}\} + 2\mathbb{P}\{\|X_{(k_{\mathrm{L}})} - x\| > r\} \\ + \sup_{z \in \mathcal{S}^{\epsilon_0}} \|\dot{\eta}(z)\| \mathbb{E}\{\|X_{(k_{\mathrm{L}})} - x\| \mathbb{1}_{\{\|X_{(k_{\mathrm{L}})} - x\| > r\}}\} \\ + \sup_{z \in \mathcal{S}^{\epsilon_0}} \|\ddot{\eta}(z)\|_{\mathrm{op}} \mathbb{E}\{\|X_{(k_{\mathrm{L}})} - x\|^2 \mathbb{1}_{\{\|X_{(k_{\mathrm{L}})} - x\| > r\}}\}. \end{aligned} \tag{27}$$

Now, by similar arguments to those leading to (18), we have that

$$\frac{\epsilon}{k_{\mathrm{L}}} \sum_{i=1}^{k_{\mathrm{L}}} \mathbb{E}(\|X_{(i)} - x\|^2 \mathbb{1}_{\{\|X_{(k_{\mathrm{L}})} - x\| \leq r\}}) = \epsilon \Big(\frac{k_{\mathrm{L}}}{n a_d \bar{f}(x)} \Big)^{2/d} \frac{d}{d+2} \{1 + o(1)\}, \tag{28}$$

uniformly for $x_0 \in \mathcal{S}_n, |t| < \epsilon_n$ and $k_{\mathrm{L}} \in K_{\beta,\tau}$. Moreover, for every $M > 0$,

$$\mathbb{P}\{\|X_{(k_{\mathrm{L}})} - x\| > r\} = q_r^n(k_{\mathrm{L}}) = O(n^{-M}), \tag{29}$$

uniformly for $x_0 \in \mathcal{S}_n, |t| < \epsilon_n$ and $k_{\mathrm{L}} \in K_{\beta,\tau}$, by (17) in Step 1. For the remaining terms, note that

$$\begin{aligned} \mathbb{E}\{\|X_{(k_{\mathrm{L}})} - x\|^2 \mathbb{1}_{\{\|X_{(k_{\mathrm{L}})} - x\| > r\}}\} &= \mathbb{P}\{\|X_{(k_{\mathrm{L}})} - x\| > r\} + \int_{r^2}^{\infty} \mathbb{P}\{\|X_{(k_{\mathrm{L}})} - x\| > \sqrt{t}\}\, dt \\ &= q_r^n(k_{\mathrm{L}}) + \int_{r^2}^{\infty} q_{\sqrt{t}}^n(k_{\mathrm{L}})\, dt. \end{aligned} \tag{30}$$

Let $t_0 = t_0(x) := 5^{2/\rho}(1 + 2^{\rho-1})^{2/\rho} \big\{ \mathbb{E}(\|X\|^\rho) + \|x\|^\rho \big\}^{2/\rho}$. Then, for $t \geq t_0$, we have

$$1 - p_{\sqrt{t}} \leq (1 + 2^{\rho-1}) \frac{\mathbb{E}(\|X\|^\rho) + \|x\|^\rho}{t^{\rho/2}} \leq \frac{1}{5}.$$

It follows by Bennett's inequality that for $\rho\{n-(n-1)^{1-\beta}\} > 4$,

$$\int_{t_0}^{\infty} q^n_{\sqrt{t}}(k_{\mathrm{L}})\,dt \le e^{k_{\mathrm{L}}}(1+2^{\rho-1})^{(n-k_{\mathrm{L}})/2}\big\{\mathbb{E}(\|X\|^\rho)+\|x\|^\rho\big\}^{(n-k_{\mathrm{L}})/2}\int_{t_0}^{\infty} t^{-\rho(n-k_{\mathrm{L}})/4}\,dt$$
$$= \frac{4e^{k_{\mathrm{L}}}5^{2/\rho}}{\rho(n-k_{\mathrm{L}})-4}(1+2^{\rho-1})^{2/\rho}\big\{\mathbb{E}(\|X\|^\rho)+\|x\|^\rho\big\}^{2/\rho}5^{-(n-k_{\mathrm{L}})/2}.$$

But, when $\beta\log(n-1) \ge (d+2)/d$ and $n \ge \max\{n_0, n_2\}$,

$$\sup_{x\in\mathcal{R}_n\cup\mathcal{S}_n^{\epsilon_n}}\|x\| \le \epsilon_0 + \left\{\frac{(n-1)^{1-\beta}c_n^d\mathbb{E}(\|X\|^\rho)}{\mu_0\beta^{d/2}\log^{d/2}(n-1)}\right\}^{1/\rho}.$$

We deduce that for every $M > 0$,

$$\sup_{k\in K_{\beta,\tau}}\sup_{x\in\mathcal{R}_n\cup\mathcal{S}_n^{\epsilon_n}}\int_{t_0}^{\infty} q^n_{\sqrt{t}}(k_{\mathrm{L}})\,dt = O(n^{-M}). \tag{31}$$

Moreover, by Bernstein's inequality, for every $M > 0$,

$$\sup_{k_{\mathrm{L}}\in K_{\beta,\tau}}\sup_{x\in\mathcal{R}_n\cup\mathcal{S}_n^{\epsilon_n}}\left\{q^n_r(k_{\mathrm{L}}) + \int_{r^2}^{t_0} q^n_{\sqrt{t}}(k_{\mathrm{L}})\,dt\right\} = O(n^{-M}). \tag{32}$$

We conclude from (15), (27), (28), (29), (30), (31) and (32), together with Jensen's inequality to deal with the third term on the right-hand side of (27), that (11) holds. With only simple modifications, we have also shown (14), which bounds R_2.

To bound R_5: Write

$$R_5 := \int_{\mathcal{S}_n} R_5(x_0)\,d\mathrm{Vol}^{d-1}(x_0)$$
$$= \int_{\mathcal{S}_n}\int_{-\epsilon_n}^{\epsilon_n} t\|\dot\psi(x_0)\|\Big[\mathbb{P}\{\hat S_n(x_0^t) < 1/2\} - \mathbb{E}\Phi\big(\hat\theta(x_0^t)\big)\Big]\,dt\,d\mathrm{Vol}^{d-1}(x_0).$$

Now by a non-uniform version of the Berry–Esseen theorem (Paditz, 1989, Theorem 1), for every $t \in (-\epsilon_n, \epsilon_n)$ and $x_0 \in \mathcal{S}_n$,

$$\big|\mathbb{P}\{\hat S_n(x_0^t) < 1/2|X^n\} - \Phi\big(\hat\theta(x_0^t)\big)\big| \le \frac{32}{k_{\mathrm{L}}(x_0^t)\hat\sigma_n(x_0^t, X^n)}\frac{1}{1+|\hat\theta(x_0^t)|^3}. \tag{33}$$

Let

$$t_n = t_n(x_0) := C\max\left\{k_{\mathrm{L}}(x_0)^{-1/2}, \left(\frac{k_{\mathrm{L}}(x_0)}{n\bar f(x_0)}\right)^{2/d} g(x_0)\right\},$$

where

$$C := \frac{4\max\{2\sup_{z\in\mathcal{S}^{\epsilon_0}}\|\dot\eta(z)\|, d\sup_{z\in\mathcal{S}^{\epsilon_0}}\|\ddot\eta(z)\|_{\mathrm{op}}\}}{(d+2)a_d^{2/d}\inf_{z\in\mathcal{S}}\|\dot\eta(z)\|}.$$

In the following we integrate the bound in (33) over the regions $|t| \le t_n$ and $|t| \in (t_n, \epsilon_n)$ separately. Define the event

$$B_{k_{\mathrm{L}}} := \Big\{ \hat{\sigma}_n(x_0^t, X^n) \ge \frac{1}{3k_{\mathrm{L}}(x_0^t)^{1/2}} \text{ for all } x_0 \in \mathcal{S}_n,\, t \in (-\epsilon_n, \epsilon_n) \Big\},$$

so that, by very similar arguments to those used to bound $\mathbb{P}(A_{k_{\mathrm{L}}}^c)$ in Step 2, we have $\mathbb{P}(B_{k_{\mathrm{L}}}^c) = O(n^{-M})$ for every $M > 0$, uniformly for $k_{\mathrm{L}} \in K_{\beta,\tau}$. It follows by (33) and Step 2 that there exists $n_4 \in \mathbb{N}$ such that for all $n \ge n_4$, $k_{\mathrm{L}} \in K_{\beta,\tau}$ and $x_0 \in \mathcal{S}_n$,

$$\begin{aligned}\Big| \int_{-t_n}^{t_n} t \Big[\mathbb{P}\{\hat{S}_n(x_0^t) < 1/2\} - \mathbb{E}\Phi\big(\hat{\theta}(x_0^t)\big) \Big]\, dt \Big| &\le \int_{-t_n}^{t_n} \mathbb{E}\bigg(\frac{32|t| \mathbb{1}_{B_{k_{\mathrm{L}}}}}{k_{\mathrm{L}}(x_0^t)\hat{\sigma}_n(x_0^t, X^n)} \bigg)\, dt + t_n^2 \mathbb{P}(B_{k_{\mathrm{L}}}^c) \\ &\le \frac{128 t_n^2}{k_{\mathrm{L}}(x_0)^{1/2}}. \end{aligned} \tag{34}$$

By Step 1, there exists $n_5 \in \mathbb{N}$ such that for $n \ge n_5$, $|t| \in (t_n, \epsilon_n)$, $x_0 \in \mathcal{S}_n$ and $k_{\mathrm{L}} \in K_{\beta,\tau}$,

$$\begin{aligned} |\mu_n(x_0^t) - 1/2| &\ge |\eta(x_0^t) - 1/2| - |\mu_n(x_0^t) - \eta(x_0^t)| \\ &\ge \frac{1}{2} \inf_{z \in \mathcal{S}} \|\dot{\eta}(z)\| |t| - \frac{1}{4} C \inf_{z \in \mathcal{S}} \|\dot{\eta}(z)\| \Big(\frac{k_{\mathrm{L}}(x_0)}{n \bar{f}(x_0)} \Big)^{2/d} g(x_0) > \frac{1}{4} \inf_{z \in \mathcal{S}} \|\dot{\eta}(z)\| |t|. \end{aligned} \tag{35}$$

Thus for $n \ge n_5$, $|t| \in (t_n, \epsilon_n)$, $x_0 \in \mathcal{S}_n$ and $k_{\mathrm{L}} \in K_{\beta,\tau}$ we have that

$$\begin{aligned} &\mathbb{P}\Big\{ |\hat{\theta}(x_0^t)| < \frac{1}{4} \inf_{z \in \mathcal{S}} \|\dot{\eta}(z)\| k_{\mathrm{L}}^{1/2}(x_0) |t| \Big\} \\ &\qquad \le \mathbb{P}\Big\{ |\hat{\mu}_n(x_0^t, X^n) - \mu_n(x_0^t)| > |\mu_n(x_0^t) - 1/2| - \frac{1}{8} \inf_{z \in \mathcal{S}} \|\dot{\eta}(z)\| |t| \Big\} \\ &\qquad \le \mathbb{P}\Big\{ |\hat{\mu}_n(x_0^t, X^n) - \mu_n(x_0^t)| > \frac{1}{8} \inf_{z \in \mathcal{S}} \|\dot{\eta}(z)\| |t| \Big\} \le \frac{64 \mathrm{Var}\{\hat{\mu}_n(x_0^t, X^n)\}}{\inf_{z \in \mathcal{S}} \|\dot{\eta}(z)\|^2 t^2}. \end{aligned} \tag{36}$$

It follows by (33), (36) and Step 3 that, for $n \ge n_5$,

$$\begin{aligned} &\Big| \int_{|t| \in (t_n, \epsilon_n)} t \Big[\mathbb{P}\{\hat{S}_n(x_0^t) < 1/2\} - \mathbb{E}\Phi\big(\hat{\theta}(x_0^t)\big) \Big]\, dt \Big| \\ &\quad \le \int_{|t| \in (t_n, \epsilon_n)} |t| \mathbb{E}\bigg(\frac{32 \mathbb{1}_{B_{k_{\mathrm{L}}}}}{k_{\mathrm{L}}(x_0^t)\hat{\sigma}_n(x_0^t, X^n)} \frac{1}{1 + \frac{1}{64} \inf_{z \in \mathcal{S}} \|\dot{\eta}(z)\|^3 k_{\mathrm{L}}(x_0)^{3/2} |t|^3} \bigg)\, dt \\ &\qquad + \int_{|t| \in (t_n, \epsilon_n)} \frac{64 \mathrm{Var}\{\hat{\mu}_n(x_0^t, X^n)\}}{\inf_{z \in \mathcal{S}} \|\dot{\eta}(z)\|^2 |t|}\, dt + \epsilon_n^2 \mathbb{P}(B_{k_{\mathrm{L}}}^c) \\ &\quad \le \frac{192}{k_{\mathrm{L}}(x_0)^{3/2}} \int_0^\infty \frac{u}{1 + \frac{1}{64} \inf_{z \in \mathcal{S}} \|\dot{\eta}(z)\|^3 u^3}\, du \\ &\qquad + \frac{128}{\inf_{z \in \mathcal{S}} \|\dot{\eta}(z)\|^2} \sup_{|t| \in (t_n, \epsilon_n)} \mathrm{Var}\{\hat{\mu}_n(x_0^t, X^n)\} \log\Big(\frac{\epsilon_n}{t_n} \Big) + \epsilon_n^2 \mathbb{P}(B_{k_{\mathrm{L}}}^c) = o\Big(\frac{1}{k_{\mathrm{L}}(x_0)} \Big) \end{aligned} \tag{37}$$

uniformly for $x_0 \in \mathcal{S}_n$ and $k_{\mathrm{L}} \in K_{\beta,\tau}$. We conclude from (34) and (37) that $|R_5| = o(\gamma_n(k_{\mathrm{L}}))$.

To bound R_6: Let $\theta(x_0^t) := -2k_{\mathrm{L}}(x_0^t)^{1/2}\{\mu_n(x_0^t) - 1/2\}$. Write

$$R_6 := \int_{\mathcal{S}_n} R_6(x_0)\, d\mathrm{Vol}^{d-1}(x_0) = R_{61} + R_{62},$$

where

$$R_{61} := \int_{\mathcal{S}_n}\int_{-\epsilon_n}^{\epsilon_n} t\|\dot{\psi}(x_0)\|\Big[\mathbb{E}\Phi\big(\hat{\theta}(x_0^t)\big) - \Phi\big(\theta(x_0^t)\big)\Big]\, dt\, d\mathrm{Vol}^{d-1}(x_0)$$

and

$$R_{62} := \int_{\mathcal{S}_n}\int_{-\epsilon_n}^{\epsilon_n} t\|\dot{\psi}(x_0)\|\Big[\Phi\big(\theta(x_0^t)\big) - \Phi\big(\bar{\theta}(x_0,t)\big)\Big]\, dt\, d\mathrm{Vol}^{d-1}(x_0).$$

To bound R_{61}: We again deal with the regions $|t| \le t_n$ and $|t| \in (t_n, \epsilon_n)$ separately. First let $\tilde{\theta}(x_0^t) := -2k_{\mathrm{L}}(x_0^t)^{1/2}\{\hat{\mu}_n(x_0^t, X^n) - 1/2\}$. Writing ϕ for the standard normal density, and using the facts that $|\hat{\theta}(x_0^t)| \ge |\tilde{\theta}(x_0^t)|$, that $\hat{\theta}(x_0^t)$ and $\tilde{\theta}(x_0^t)$ have the same sign, and that $|x\phi(x)| \le 1$, we have

$$\begin{aligned}
&\left|\int_{-t_n}^{t_n} t\big[\mathbb{E}\Phi\big(\hat{\theta}(x_0^t)\big) - \Phi\big(\theta(x_0^t)\big)\big]\, dt\right| \\
&\le \int_{-t_n}^{t_n} |t|\mathbb{E}\Big\{|\hat{\theta}(x_0^t) - \tilde{\theta}(x_0^t)|\phi\big(\tilde{\theta}(x_0^t)\big)\mathbb{1}_{A_{k_{\mathrm{L}}}} + |\tilde{\theta}(x_0^t) - \theta(x_0^t)|\Big\}\, dt + t_n^2\mathbb{P}(A_{k_{\mathrm{L}}}^c) \\
&\le \int_{-t_n}^{t_n} |t|\left[\mathbb{E}\left\{\mathbb{1}_{A_{k_{\mathrm{L}}}}\left|\frac{1}{2k_{\mathrm{L}}(x_0^t)^{1/2}\hat{\sigma}_n(x_0^t, X^n)} - 1\right|\right\} + 2k_{\mathrm{L}}(x_0^t)^{1/2}\mathrm{Var}^{1/2}\{\hat{\mu}_n(x_0^t, X^n)\}\right] dt \\
&\qquad\qquad + t_n^2\mathbb{P}(A_{k_{\mathrm{L}}}^c) = o(t_n^2)
\end{aligned}$$

uniformly for $x_0 \in \mathcal{S}_n$ and $k_{\mathrm{L}} \in K_{\beta,\tau}$. Note that for $|t| \in (t_n, \epsilon_n)$ and $x_0 \in \mathcal{S}_n$, we have when $\epsilon_n < \epsilon_0$ and $n \ge n_5$ that

$$\begin{aligned}
&\mathbb{E}\big\{\mathbb{1}_{A_{k_{\mathrm{L}}}\cap B_{k_{\mathrm{L}}}}\big|\hat{\theta}(x_0^t) - \theta(x_0^t)\big|\big\} \\
&\le \mathbb{E}\left\{\frac{\mathbb{1}_{A_{k_{\mathrm{L}}}\cap B_{k_{\mathrm{L}}}}}{\hat{\sigma}_n(x_0^t, X^n)}|\hat{\mu}_n(x_0^t, X^n) - \mu_n(x_0^t)| + \mathbb{1}_{A_{k_{\mathrm{L}}}\cap B_{k_{\mathrm{L}}}}|\theta(x_0^t)|\left|\frac{1}{2k_{\mathrm{L}}(x_0^t)^{1/2}\hat{\sigma}_n(x_0^t, X^n)} - 1\right|\right\} \\
&\le 3k_{\mathrm{L}}(x_0)^{1/2}\mathrm{Var}^{1/2}\{\hat{\mu}_n(x_0^t, X^n)\} \\
&\qquad + \frac{5}{2}k_{\mathrm{L}}(x_0)^{1/2}\sup_{z\in\mathcal{S}^{\epsilon_0}}\|\dot{\eta}(z)\||t|\mathbb{E}\left\{\mathbb{1}_{A_{k_{\mathrm{L}}}\cap B_{k_{\mathrm{L}}}}\left|\frac{1}{2k_{\mathrm{L}}(x_0^t)^{1/2}\hat{\sigma}_n(x_0^t, X^n)} - 1\right|\right\}.
\end{aligned} \tag{38}$$

Thus by (35), (36), (38) and Step 3, for $\epsilon_n < \epsilon_0$ and $n \geq n_5$,

$$\begin{aligned}
&\int_{|t|\in(t_n,\epsilon_n)} |t|\big|\mathbb{E}\Phi\big(\hat{\theta}(x_0^t)\big) - \Phi\big(\theta(x_0^t)\big)\big|\,dt \\
&\leq \int_{|t|\in(t_n,\epsilon_n)} |t|\mathbb{E}\Big\{\mathbb{1}_{A_{k_{\mathrm{L}}}\cap B_{k_{\mathrm{L}}}}\big|\hat{\theta}(x_0^t) - \theta(x_0^t)\big|\Big\}\phi\Big(\frac{1}{4}\inf_{z\in\mathcal{S}}\|\dot{\eta}(z)\|k_{\mathrm{L}}^{1/2}(x_0)|t|\Big)\,dt + \mathbb{P}(A_{k_{\mathrm{L}}}^c \cup B_{k_{\mathrm{L}}}^c) \\
&+ \frac{128}{\inf_{z\in\mathcal{S}}\|\dot{\eta}(z)\|^2}\sup_{|t|\in(t_n,\epsilon_n)}\mathrm{Var}\{\hat{\mu}_n(x_0^t, X^n)\}\log\Big(\frac{\epsilon_n}{t_n}\Big) = o\Big(\frac{1}{k_{\mathrm{L}}(x_0)}\Big) \qquad (39)
\end{aligned}$$

uniformly for $x_0 \in \mathcal{S}_n$ and $k_{\mathrm{L}} \in K_{\beta,\tau}$.

To bound R_{62}: Let

$$u(x) := k_{\mathrm{L}}(x)^{1/2}\Big(\frac{k_{\mathrm{L}}(x)}{n\bar{f}(x)}\Big)^{2/d}.$$

Given $\epsilon > 0$ small enough that $\epsilon^2 + \frac{\epsilon}{2\inf_{x\in\mathcal{S}}\|\dot{\eta}(x_0)\|} < 1/2$, by Step 1 there exists $n_6 \in \mathbb{N}$ such that for $n \geq n_6$, $k_{\mathrm{L}} \in K_{\beta,\tau}$, $x_0 \in \mathcal{S}_n$ and $|t| < \epsilon_n$,

$$\big|\theta(x_0^t) - \bar{\theta}(x_0,t)\big| \leq \epsilon^2\big\{|t|k_{\mathrm{L}}(x_0)^{1/2} + u(x_0)g(x_0)\big\}.$$

By decreasing ϵ and increasing n_6 if necessary, it follows that

$$\big|\Phi\big(\theta(x_0^t)\big) - \Phi\big(\bar{\theta}(x_0,t)\big)\big| \leq \epsilon^2\big\{|t|k_{\mathrm{L}}(x_0)^{1/2} + u(x_0)g(x_0)\big\}\phi\Big(\frac{1}{2}\bar{\theta}(x_0,t)\Big),$$

for all $n \geq n_6$, $k_{\mathrm{L}} \in K_{\beta,\tau}$, and $x_0 \in \mathcal{S}_n$, $t \in (-\epsilon_n, \epsilon_n)$ satisfying $2\epsilon u(x_0)g(x_0)\|\dot{\eta}(x_0)\| \leq |\bar{\theta}(x_0,t)|$. Substituting $u = \bar{\theta}(x_0,t)/2$, it follows that there exists $C^* > 0$ such that for all $n \geq n_6$ and all $k_{\mathrm{L}} \in K_{\beta,\tau}$,

$$\begin{aligned}
|R_{62}| &\leq \int_{\mathcal{S}_n}\int_{|u|\leq\epsilon u(x_0)g(x_0)\|\dot{\eta}(x_0)\|}\frac{2\bar{f}(x_0)}{\|\dot{\eta}(x_0)\|k_{\mathrm{L}}(x_0)}|u + u(x_0)a(x_0)|\,du\,d\mathrm{Vol}^{d-1}(x_0) \\
&+ \int_{\mathcal{S}_n}\int_{-\infty}^{\infty}\frac{2\bar{f}(x_0)|u + u(x_0)a(x_0)|}{\|\dot{\eta}(x_0)\|^2 k_{\mathrm{L}}(x_0)}\big\{\epsilon^2|u + u(x_0)a(x_0)| + \epsilon|u|\big\}\phi(u)\,du\,d\mathrm{Vol}^{d-1}(x_0) \\
&\leq C^*\epsilon\gamma_n(k_{\mathrm{L}}). \qquad (40)
\end{aligned}$$

The combination of (39) and (40) yields the desired error bound on $|R_6|$ in (26) and therefore completes the proof. □

6.2 Proof of Theorem 1

Proof of Theorem 1. Let $k \in K_\beta$, and note that since $k_{\mathrm{L}}(x) = k$ is constant, we have that $c_n = \sup_{x_0\in\mathcal{S}:\bar{f}(x_0)\geq k/(n-1)} g(x_0)$, and $\delta_n = \frac{k}{n-1}c_n^d\log^d(\frac{n-1}{k})$. Now let

$$\mathcal{R}_n = \{x \in \mathbb{R}^d : \bar{f}(x) > \delta_n\} \cap \mathcal{X}_{\bar{f}},$$

and let $n_0 \in \mathbb{N}$ be large enough that $\mathcal{R}_n$ is non-empty for $n \geq n_0$, so that, by Assumption **(A.1)**, for $n \geq n_0$ it is an open subset of $\mathbb{R}^d$, and therefore a d-dimensional manifold. For $n \geq n_0$, we may apply Theorem 4 with $k_{\mathrm{L}}(x) = k$ for all $x \in \mathbb{R}^d$ to deduce that

$$R_{\mathcal{R}_n}(\hat{C}_n^{k\mathrm{nn}}) - R_{\mathcal{R}_n}(C^{\mathrm{Bayes}}) = B_{1,n}\frac{1}{k} + B_{2,n}\Big(\frac{k}{n}\Big)^{4/d} + o(\gamma_n(k)) + O\big\{P_X\big((\partial\mathcal{R}_n)^{\epsilon_n} \cap \mathcal{S}^{\epsilon_n}\big)\big\}$$

uniformly for $k \in K_\beta$, where

$$B_{1,n} := \int_{\mathcal{S}_n} \frac{\bar{f}(x_0)}{4\|\dot{\eta}(x_0)\|}\, d\mathrm{Vol}^{d-1}(x_0) \quad \text{and} \quad B_{2,n} := \int_{\mathcal{S}_n} \frac{\bar{f}(x_0)^{1-4/d}}{\|\dot{\eta}(x_0)\|} a(x_0)^2\, d\mathrm{Vol}^{d-1}(x_0),$$

and $\mathcal{S}_n := \mathcal{S} \cap \mathcal{R}_n$. We now show that, under the conditions of part (i), $B_{1,n}$ and $B_{2,n}$ are well approximated by integrals over the whole of the manifold $\mathcal{S}$, and that these integrals are finite. First, by Assumptions **(A.3)** and **(A.4)**(ρ),

$$B_1 = \int_{\mathcal{S}} \frac{\bar{f}(x_0)}{4\|\dot{\eta}(x_0)\|}\, d\mathrm{Vol}^{d-1}(x_0) \leq \frac{1}{4}\frac{1}{\inf_{x_0\in\mathcal{S}} \|\dot{\eta}(x_0)\|} \int_{\mathcal{S}} \bar{f}(x_0)\, d\mathrm{Vol}^{d-1}(x_0) < \infty.$$

Moreover,

$$\begin{aligned} B_1 - B_{1,n} &= \int_{\mathcal{S}\setminus\mathcal{R}_n} \frac{\bar{f}(x_0)}{4\|\dot{\eta}(x_0)\|}\, d\mathrm{Vol}^{d-1}(x_0) \\ &\leq \frac{1}{4}\frac{1}{\inf_{x_0\in\mathcal{S}} \|\dot{\eta}(x_0)\|} \int_{\mathcal{S}\setminus\mathcal{R}_n} \bar{f}(x_0)\, d\mathrm{Vol}^{d-1}(x_0) \to 0, \end{aligned}$$

uniformly for $k \in K_\beta$. By Assumptions **(A.2)**, **(A.3)** and **(A.4)**(ρ) and the fact that $\rho/(\rho+d) > 4/d$, we have that

$$\begin{aligned} B_2 &= \int_{\mathcal{S}} \frac{\bar{f}(x_0)^{1-4/d}}{\|\dot{\eta}(x_0)\|} a(x_0)^2\, d\mathrm{Vol}^{d-1}(x_0) \\ &\leq \sup_{x_0\in\mathcal{S}}\Big\{\frac{a(x_0)^2 \bar{f}(x_0)^{\rho/(\rho+d)-4/d}}{\|\dot{\eta}(x_0)\|}\Big\} \int_{\mathcal{S}} \bar{f}(x_0)^{d/(\rho+d)}\, d\mathrm{Vol}^{d-1}(x_0) < \infty. \end{aligned}$$

Similarly,

$$\begin{aligned} B_2 - B_{2,n} &= \int_{\mathcal{S}\setminus\mathcal{R}_n} \frac{\bar{f}(x_0)^{1-4/d}}{\|\dot{\eta}(x_0)\|} a(x_0)^2\, d\mathrm{Vol}^{d-1}(x_0) \\ &\leq \sup_{x_0\in\mathcal{S}}\Big\{\frac{a(x_0)^2 \bar{f}(x_0)^{\rho/(\rho+d)-4/d}}{\|\dot{\eta}(x_0)\|}\Big\} \int_{\mathcal{S}\setminus\mathcal{R}_n} \bar{f}(x_0)^{d/(\rho+d)}\, d\mathrm{Vol}^{d-1}(x) \to 0, \end{aligned}$$

uniformly for $k \in K_\beta$, as $n \to \infty$. A similar argument shows that $\gamma_n(k) = O(1/k + (k/n)^{4/d})$, uniformly for $k \in K_\beta$.

Finally, we bound $P_X\big((\partial\mathcal{R}_n)^{\epsilon_n} \cap \mathcal{S}^{\epsilon_n}\big)$ and $R_{\mathcal{R}_n^c}(\hat{C}_n^{\mathrm{knn}}) - R_{\mathcal{R}_n^c}(C^{\mathrm{Bayes}})$. Suppose that $x \in (\partial\mathcal{R}_n)^{\epsilon_n} \cap \mathcal{S}^{\epsilon_n}$. Then there exists $z \in \partial\mathcal{R}_n \cap B_{\epsilon_n}(x) \cap \mathcal{S}^{2\epsilon_n}$ with $\bar{f}(z) = \delta_n$. By Assumption **(A.2)** we have that

$$\Big|\frac{\bar{f}(x)}{\bar{f}(z)} - 1\Big| \leq g(z)\|x-z\| + \frac{1}{2}g(z)\|x-z\|^2 \leq \frac{1+\epsilon_n/2}{\beta^{1/2}\log^{1/2}(n-1)}. \tag{41}$$

Thus there exists $n_1 \in \mathbb{N}$ such that $(\partial\mathcal{R}_n)^{\epsilon_n} \cap \mathcal{S}^{\epsilon_n} \subseteq \{x \in \mathbb{R}^d : \bar{f}(x) \leq 2\delta_n\}$ for $n \geq n_1$. By the moment assumption in **(A.4)**(ρ) and Hölder's inequality, observe that for any $\alpha \in (0,1)$, $n \geq n_1$ and $\epsilon > 0$,

$$\begin{aligned}
P_X\big((\partial\mathcal{R}_n)^{\epsilon_n} \cap \mathcal{S}^{\epsilon_n}\big) &\leq \mathbb{P}\{\bar{f}(X) \leq 2\delta_n\} \leq (2\delta_n)^{\frac{\rho(1-\alpha)}{\rho+d}} \int_{x:\bar{f}(x)\leq 2\delta_n} \bar{f}(x)^{1-\frac{\rho(1-\alpha)}{\rho+d}}\,dx \\
&\leq (2\delta_n)^{\frac{\rho(1-\alpha)}{\rho+d}} \Big\{\int_{\mathbb{R}^d} (1+\|x\|^\rho)\bar{f}(x)\,dx\Big\}^{1-\frac{\rho(1-\alpha)}{\rho+d}} \Big\{\int_{\mathbb{R}^d} \frac{1}{(1+\|x\|^\rho)^{\frac{d+\rho\alpha}{\rho(1-\alpha)}}}\,dx\Big\}^{\frac{\rho(1-\alpha)}{\rho+d}} \\
&= o\bigg(\Big(\frac{k}{n}\Big)^{\frac{\rho(1-\alpha)}{\rho+d}-\epsilon}\bigg),
\end{aligned} \tag{42}$$

uniformly for $k \in K_\beta$. Moreover,

$$R_{\mathcal{R}_n^c}(\hat{C}_n^{\mathrm{knn}}) - R_{\mathcal{R}_n^c}(C^{\mathrm{Bayes}}) \leq P_X(\mathcal{R}_n^c) \leq \mathbb{P}\{\bar{f}(X) \leq 2\delta_n\},$$

so the same bound (42) applies for this region. Since $\rho/(\rho+d) > 4/d$, this completes the proof of part (i).

For part (ii), in contrast to part (i), the dominant contribution to the excess risk could now arise from the tail of the distribution. First, as in part (i), we have $B_{1,n} \to B_1 < \infty$, uniformly for $k \in K_\beta$. Furthermore, using Assumptions **(A.3)** and **(A.4)**(ρ) and the fact that $4/d > \rho/(\rho+d)$, we see that

$$\begin{aligned}
B_{2,n}\Big(\frac{k}{n}\Big)^{4/d} &\leq \delta_n^{\rho/(\rho+d)} \int_{\mathcal{S}_n} \frac{\delta_n^{4/d-\rho/(\rho+d)}}{c_n^4\log^4((n-1)/k)} \frac{\bar{f}(x_0)^{1-4/d}}{\|\dot{\eta}(x_0)\|} a(x_0)^2\,d\mathrm{Vol}^{d-1}(x_0) \\
&\leq \sup_{x_0\in\mathcal{S}_n}\Big\{\frac{a(x_0)^2}{\|\dot{\eta}(x_0)\|}\Big\} \frac{\delta_n^{\rho/(\rho+d)}}{c_n^4\log^4((n-1)/k)} \int_{\mathcal{S}} \bar{f}(x_0)^{d/(\rho+d)}\,d\mathrm{Vol}^{d-1}(x_0) \\
&= o((k/n)^{\rho/(\rho+d)-\epsilon}),
\end{aligned}$$

for every $\epsilon > 0$, uniformly for $k \in K_\beta$, where the final equality follows from the fact that $\sup_{x_0\in\mathcal{S}_n} a^2(x_0)/c_n^2$ is bounded. We can also bound $\gamma_n(k)$ by the same argument, so the result follows in the same way as in part (i). □

6.3 Proof of claim in Example 1

Proof of claim in Example 1. Fix $\epsilon > 0$ and $k \in K_\beta$, let

$$\mathcal{T}_n := (0, 1/2) \times \big((1+\epsilon)\log(n/k), \infty\big),$$

and for $\gamma > 0$, let

$$B_{k,\gamma} = \bigcap_{x=(x_1,x_2)\in\mathcal{T}_n} \{\gamma < \|X_{(k+1)}(x) - x\| < x_2 - 1\}.$$

Now, for $\epsilon\beta \log n > 4$ and $\gamma \in [2, \epsilon\log(n/k)/2)$,

$$\mathbb{P}(B_{k,\gamma}^c) \le \mathbb{P}(T \ge k+1) + \mathbb{P}(T' \le k),$$

where $T \sim \mathrm{Bin}(n, p_\gamma^*)$, $T' \sim \mathrm{Bin}(n, p_*)$,

$$p_\gamma^* := \int_0^1 \int_{(1+\epsilon)\log(n/k)-\gamma}^{\infty} t_1 \exp(-t_2)\, dt_1 dt_2 \le \frac{1}{2}\Big(\frac{k}{n}\Big)^{1+\epsilon} e^\gamma \le \frac{1}{2}\Big(\frac{k}{n}\Big)^{1+\epsilon/2},$$

$$p_* := \int_0^1 \int_{3-3^{1/2}}^{3+3^{1/2}} t_1 \exp(-t_2)\, dt_1 dt_2 \ge \frac{1}{8}.$$

Therefore, there exists $n_0 \in \mathbb{N}$ such that $np_* - (k+1) \ge k/2$ and $k + 1 - np_\gamma^* \ge k/2$ for all $k \in K_\beta$, $\gamma \in [2, \epsilon\log(n/k)/2)$ and $n \ge n_0$. It follows by Bernstein's inequality that $\sup_{k\in K_\beta} \sup_{\gamma\in[2,\epsilon\log(n/k)/2)} \mathbb{P}(B_{k,\gamma}^c) = O(n^{-M})$ for every $M > 0$.

Now, for $x = (x_1, x_2) \in \mathcal{T}_n$, $\epsilon\beta\log n > 4$ and $\gamma \in [2, x_2 - 1)$, we have that

$$\frac{\int_{B_\gamma(x)} \eta(t)\bar{f}(t)\, dt}{\int_{B_\gamma(x)} \bar{f}(t)\, dt} = \frac{\int_0^1 \int_{x_2-\{\gamma^2-(t_1-x_1)^2\}^{1/2}}^{x_2+\{\gamma^2-(t_1-x_1)^2\}^{1/2}} t_1^2 e^{-t_2}\, dt_2\, dt_1}{\int_0^1 \int_{x_2-\{\gamma^2-(t_1-x_1)^2\}^{1/2}}^{x_2+\{\gamma^2-(t_1-x_1)^2\}^{1/2}} t_1 e^{-t_2}\, dt_2\, dt_1} = \frac{\int_0^1 t_1^2 \sinh(\{\gamma^2-(t_1-x_1)^2\}^{1/2})\, dt_1}{\int_0^1 t_1 \sinh(\{\gamma^2-(t_1-x_1)^2\}^{1/2})\, dt_1}$$
$$\ge \frac{2}{3}\frac{\sinh\big((\gamma^2-1)^{1/2}\big)}{\sinh(\gamma)} \ge \frac{2}{3}\frac{\sinh(3^{1/2})}{\sinh(2)} > \frac{1}{2}.$$

Our next observation is that for $\gamma \in [0, \infty)$ and $x_{(k+1)} \in \mathbb{R}^d$ such that $\|x_{(k+1)} - x\| = \gamma$, we have that $(X_{(1)}, Y_{(1)}, \ldots, X_{(k)}, Y_{(k)}) | (X_{(k+1)} = x_{(k+1)}) \stackrel{d}{=} (\tilde{X}_{(1)}, \tilde{Y}_{(1)}, \ldots, \tilde{X}_{(k)}, \tilde{Y}_{(k)})$, where the pairs $(\tilde{X}_1, \tilde{Y}_1), \ldots, (\tilde{X}_k, \tilde{Y}_k)$ are independent and identically distributed, and then $(\tilde{X}_{(1)}, \tilde{Y}_{(1)}), \ldots, (\tilde{X}_{(k)}, \tilde{Y}_{(k)})$ is a reordering such that $\|\tilde{X}_{(1)} - x\| \le \ldots \le \|\tilde{X}_{(k)} - x\|$. Here $\tilde{X}_1 \stackrel{d}{=} X | (\|X - x\| \le \gamma)$ and $\mathbb{P}(\tilde{Y}_1 = 1 | \tilde{X}_1 = x) = \eta(x)$. Writing $\tilde{S}_n(x) := \frac{1}{k}\sum_{i=1}^k \mathbb{1}_{\{\tilde{Y}_i = 1\}}$ we therefore have by Hoeffding's inequality that, for $x \in \mathcal{T}_n$, $\epsilon\beta\log n > 4$ and $\|x_{(k+1)} - x\| \in [2, x_2 - 1)$,

$$\begin{aligned}
\mathbb{P}\{\hat{S}_n(x) < 1/2 \big| X_{(k+1)} = x_{(k+1)}\} &= \mathbb{P}\{\tilde{S}_n(x) < 1/2\} \\
&= \mathbb{P}\{\tilde{S}_n(x) - \mathbb{E}\tilde{S}_n(x) < -(\mathbb{E}\eta(\tilde{X}_1) - 1/2)\} \\
&\le \exp\left(-2k\Big(\frac{2}{3}\frac{\sinh(3^{1/2})}{\sinh(2)} - \frac{1}{2}\Big)^2\right) = O(n^{-M})
\end{aligned}$$

for all $M > 0$, uniformly for $k \in K_\beta$. Writing $P_{(k+1)}$ for the marginal distribution of $X_{(k+1)}$, we deduce that

$$\begin{aligned}\mathbb{P}\{\hat{S}_n(x) < 1/2\} &\le \mathbb{P}\{\hat{S}_n(x) < 1/2, \|X_{(k+1)} - x\| \in [2, x_2 - 1)\} + \mathbb{P}(B_{k,2}^c) \\ &= \int_{B_{x_2-1}(x)\setminus B_2(x)} \mathbb{P}\{\hat{S}_n(x) < 1/2 \big| X_{(k+1)} = x_{(k+1)}\}\, dP_{(k+1)}(x_{(k+1)}) + O(n^{-M}) \\ &= O(n^{-M})\end{aligned}$$

for all $M > 0$, uniformly for $k \in K_\beta$. We conclude that for every $M > 0$,

$$\begin{aligned}R_{\mathcal{T}_n}(\hat{C}_n^{k\mathrm{nn}}) - R_{\mathcal{T}_n}(C^{\mathrm{Bayes}}) &= \int_{\mathcal{T}_n} \Big[\mathbb{P}\{\hat{S}_n(x) < 1/2\} - \mathbb{1}_{\{\eta(x)<1/2\}}\Big]\{2\eta(x) - 1\}\bar{f}(x)\, dx \\ &= \int_{(1+\epsilon)\log(n/k)}^{\infty} \int_0^{1/2} \mathbb{P}\{\hat{S}_n(x) \ge 1/2\}(1 - 2x_1)x_1 \exp(-x_2)\, dx_1\, dx_2 \\ &= \frac{1}{24}\Big(\frac{k}{n}\Big)^{1+\epsilon} + O(n^{-M}),\end{aligned}$$

uniformly for $k \in K_\beta$, which establishes the claim (4). □

6.4 Proofs of results from Section 4

Proof of Theorem 2. Recall that

$$k_{\mathrm{O}}(x) = \max\big[\lceil (n-1)^\beta \rceil, \min\big\{\big\lfloor B\big\{\bar{f}(x)(n-1)\big\}^{4/(d+4)}\big\rfloor, \lfloor (n-1)^{1-\beta}\rfloor\big\}\big],$$

and define

$$\delta_{n,\mathrm{O}}(x) := \frac{k_{\mathrm{O}}(x)}{n-1} c_n^d \log^d\Big(\frac{n-1}{k_{\mathrm{O}}(x)}\Big),$$

where $c_n := \sup_{x_0 \in \mathcal{S}: \bar{f}(x_0) \ge k_{\mathrm{O}}(x_0)/(n-1)} g(x_0)$. For $\alpha \in ((1 + d/4)\beta, 1)$ let

$$\mathcal{R}_n = \{x \in \mathbb{R}^d : \bar{f}(x) > (n-1)^{-(1-\alpha)}\} \cap \mathcal{X}_{\bar{f}}.$$

Then there exists $n_0 \in \mathbb{N}$ such that for $n \ge n_0$ we have $\mathcal{R}_n \subseteq \big\{x \in \mathbb{R}^d : \bar{f}(x) \ge \delta_{n,\mathrm{O}}(x)\big\}$ and by Assumption **(A.1)** we then have that $\mathcal{R}_n$ is a d-dimensional manifold. There exists $n_1 \in \mathbb{N}$ such that for all $n \ge n_1$ and $x \in \mathcal{R}_n \cap \mathcal{S}^{\epsilon_0}$ we have that $k_{\mathrm{O}}(x) = \big\lfloor B\big\{\bar{f}(x)(n-1)\big\}^{4/(d+4)}\big\rfloor$. By **(A.2)**, we therefore have that $k_{\mathrm{O}} \in K_{\beta,\tau}$ for some $\tau = \tau_n$ with $\tau_n \searrow 0$. We deduce from Theorem 4 that

$$R(\hat{C}_n^{k_{\mathrm{O}}\mathrm{nn}}) - R(C^{\mathrm{Bayes}}) = B_{3,n} n^{-4/(d+4)} + o(\gamma_n(k_{\mathrm{O}})) + O\big\{P_X\big((\partial\mathcal{R}_n)^{\epsilon_n} \cap \mathcal{S}^{\epsilon_n}\big) + P_X(\mathcal{R}_n^c)\big\}$$

as $n \to \infty$, where

$$B_{3,n} := \int_{\mathcal{S}_n} \frac{\bar{f}(x_0)^{d/(d+4)}}{\|\dot{\eta}(x_0)\|}\Big\{\frac{1}{4B} + B^{4/d} a(x_0)^2\Big\}\, d\mathrm{Vol}^{d-1}(x_0).$$

By a similar argument to that in (41) we have that if $x \in (\partial\mathcal{R}_n)^{\epsilon_n} \cap \mathcal{S}^{\epsilon_n}$ then $\bar{f}(x) \leq 2(n-1)^{-(1-\alpha)}$. But, by Markov's inequality and Hölder's inequality, for $\tilde{\alpha} \in (0,1)$,

$$\begin{aligned}\mathbb{P}\{\bar{f}(X) \leq 2(n-1)^{-(1-\alpha)}\} &\leq \{2(n-1)^{-(1-\alpha)}\}^{\frac{\rho(1-\tilde{\alpha})}{\rho+d}} \int_{\mathbb{R}^d} \bar{f}(x)^{1-\frac{\rho(1-\tilde{\alpha})}{\rho+d}}\,dx \\ &\leq \{2(n-1)^{-(1-\alpha)}\}^{\frac{\rho(1-\tilde{\alpha})}{\rho+d}} \Big\{\int_{\mathbb{R}^d} (1+\|x\|^\rho)\bar{f}(x)\,dx\Big\}^{1-\frac{\rho(1-\tilde{\alpha})}{\rho+d}} \\ &\quad \Big\{\int_{\mathbb{R}^d} \frac{1}{(1+\|x\|^\rho)^{(\rho+d)/\{\rho(1-\tilde{\alpha})\}-1}}\,dx\Big\}^{\frac{\rho(1-\tilde{\alpha})}{\rho+d}}. \qquad (43)\end{aligned}$$

Thus, if $\rho > 4$, then we can choose $\alpha \in ((1+d/4)\beta, d(\rho-4)/\{\rho(d+4)\})$ and $\tilde{\alpha} < 1 - 4(\rho+d)/\{\rho(1-\alpha)(d+4)\}$ in (43) to conclude that

$$P_X(\mathcal{R}_n^c) \leq \mathbb{P}\{\bar{f}(X) \leq 2(n-1)^{-(1-\alpha)}\} = o(n^{-4/(d+4)}).$$

Moreover, by very similar arguments to those given in the proof of Theorem 1, $\gamma_n(k_{\mathrm{O}}) = O(n^{-4/(d+4)})$ and $B_{3,n} \to B_3$ as $n \to \infty$. This concludes the proof of part (i).

On the other hand, if $\rho \leq 4$, then choosing both $\tilde{\alpha} > 0$ and $\alpha > (1+d/4)\beta$ to be sufficiently small, we find from (43) that

$$B_{3,n}n^{-4/(d+4)} + \gamma_n(k_{\mathrm{O}}) + P_X\big((\partial\mathcal{R}_n)^{\epsilon_n} \cap \mathcal{S}^{\epsilon_n}\big) + P_X(\mathcal{R}_n^c) = o\Big(\Big(\frac{1}{n}\Big)^{\frac{\rho}{\rho+d}-\beta-\epsilon}\Big),$$

for every $\epsilon > 0$. This proves part (ii). □

Proof of Theorem 3. We prove parts (i) and (ii) of the theorem simultaneously, by appealing to the corresponding arguments in the proof of Theorem 2. First, as in the proof of Theorem 2, for $\alpha \in \big((1+d/4)\beta, 1\big)$, we define $\mathcal{R}_n = \{x \in \mathbb{R}^d : \bar{f}(x) > (n-1)^{-(1-\alpha)}\} \cap \mathcal{X}_{\bar{f}}$ and introduce the following class of functions: for $\tau > 0$, let

$$\mathcal{H}_{n,\tau} := \Big\{h : \mathbb{R}^d \to \mathbb{R} : h \text{ continuous}, \sup_{x\in\mathcal{R}_n}\Big|\frac{\bar{f}(x)}{h(x)} - 1\Big| \leq \tau\Big\}.$$

Let $\tau = \tau_n := 2(n-1)^{-\alpha/2}$. We first show that $\hat{f}_m \in \mathcal{H}_{n,\tau}$ with high probability. For $x \in \mathcal{R}_n$,

$$\Big|\frac{\hat{f}_m(x)}{\bar{f}(x)} - 1\Big| \leq (n-1)^{1-\alpha}|\hat{f}_m(x) - \bar{f}(x)| \leq (n-1)^{1-\alpha}\|\hat{f}_m - \bar{f}\|_\infty.$$

Now

$$\|\hat{f}_m - \bar{f}\|_\infty \leq \|\hat{f}_m - \mathbb{E}\hat{f}_m\|_\infty + \|\mathbb{E}\hat{f}_m - \bar{f}\|_\infty. \qquad (44)$$

To bound the first term in (44), by Giné and Guillou (2002, Theorem 2.1), there exist $C, L > 0$, such that

$$\mathbb{P}(\|\hat{f}_m - \mathbb{E}\hat{f}_m\|_\infty \geq sm^{-\gamma/(d+2\gamma)}) \leq L\exp\Big(\frac{-\log(1+C/(4L))A^d s^2}{LC\|\bar{f}\|_\infty R(K)}\Big), \qquad (45)$$

for all $CA^{-d/2}\|\bar{f}\|_\infty^{1/2}R(K)^{1/2}\log^{1/2}\big(\frac{\|K\|_\infty m^{d/(2(d+2\gamma))}}{\|\bar{f}\|_\infty^{1/2}A^{d/2}R(K)^{1/2}}\big) \leq s \leq \frac{C\|\bar{f}\|_\infty R(K)m^{\gamma/(d+2\gamma)}}{\|K\|_\infty}$. Then, by applying the bound in (45) with $s = s_0 := (n-1)^{\alpha/2}m_0^{\gamma/(d+2\gamma)}$, since $m \geq m_0(n-1)^{d/\gamma+2}$, we have that, for large n,

$$\mathbb{P}\Big\{\|\hat{f}_m - \mathbb{E}\hat{f}_m\|_\infty \geq \frac{1}{(n-1)^{1-\alpha/2}}\Big\} = \mathbb{P}\Big\{\|\hat{f}_m - \mathbb{E}\hat{f}_m\|_\infty \geq s_0 m^{-\gamma/(d+2\gamma)}\Big\}$$
$$\leq L\exp\Big(\frac{-\log(1+C/(4L))A^d(n-1)^\alpha m_0^{\gamma/(d+2\gamma)}}{LC\|\bar{f}\|_\infty R(K)}\Big) = O(n^{-M}),$$

for all $M > 0$. For the second term in (44), by a Taylor expansion and **(A.5)(γ)**, we have that, for all n sufficiently large,

$$\|\mathbb{E}\hat{f}_m - \bar{f}\|_\infty \leq \lambda A^\gamma m^{-\gamma/(d+2\gamma)}\int_{\mathbb{R}^d}\|z\|^\gamma|K(z)|\,dz = \frac{\lambda A^\gamma m_0^{-\gamma/(d+2\gamma)}}{n-1}\int_{\mathbb{R}^d}\|z\|^\gamma|K(z)|\,dz.$$

It follows that $\mathbb{P}(\hat{f}_m \notin \mathcal{H}_{n,\tau}) = O(n^{-M})$ for all $M > 0$, with $\tau = 2(n-1)^{-\alpha/2}$.

Now, for $h \in \mathcal{H}_{n,\tau}$, let

$$k_h(x) := \max\Big[\lceil(n-1)^\beta\rceil, \min\big\{\lfloor B\{h(x)(n-1)\}^{4/(d+4)}\rfloor, \lfloor(n-1)^{1-\beta}\rfloor\big\}\Big].$$

Let $c_n := \sup_{x_0\in\mathcal{S}:\bar{f}(x_0)\geq k_h(x_0)/(n-1)} g(x_0)$, and let

$$\delta_{n,h}(x) := \frac{k_h(x)}{n-1}c_n^d\log^d\Big(\frac{n-1}{k_h(x)}\Big).$$

Then there exists $n_0 \in \mathbb{N}$ such that for $n \geq n_0$ and $h \in \mathcal{H}_{n,\tau}$, we have $\mathcal{R}_n \subseteq \big\{x \in \mathbb{R}^d : \bar{f}(x) \geq \delta_{n,h}(x)\big\}$ and $k_h \in K_{\beta,\tau}$. We can therefore apply Theorem 4 (similarly to the application in the proof of Theorem 2) to conclude that for every $M > 0$,

$$R(\hat{C}_n^{k_h\mathrm{nn}}) - R(C^{\mathrm{Bayes}}) = B_{3,n}n^{-4/(d+4)}\{1+o(1)\} + o(\gamma_n(k_h))$$
$$+ O\big\{P_X\big((\partial\mathcal{R}_n)^{\epsilon_n}\cap\mathcal{S}^{\epsilon_n}\big) + P_X(\mathcal{R}_n^c)\big\} + O(n^{-M}),$$

uniformly for $h \in \mathcal{H}_{n,\tau}$, where $B_{3,n}$ was defined in the proof of Theorem 2. The proof of both parts (i) and (ii) is now completed by following the relevant steps in the proof of Theorem 2. □

Acknowledgements

The research of the third author is supported by an Engineering and Physical Sciences Research Council Fellowship and a grant from the Leverhulme Trust.

References

Audibert, J.-Y. and Tsybakov, A. B. (2007). Fast learning rates for plug-in classifiers. *Ann. Statist.*, **35**, 608–633.

Berrett, T. B., Samworth, R. J. and Yuan, M. (2016). Efficient multivariate entropy estimation via k-nearest neighbour distances. *ArXiv e-prints*, 1606.00304.

Biau, G., Cérou, F. and Guyader, A. (2010). On the rate of convergence of the bagged nearest neighbor estimate. *J. Mach. Learn. Res.*, **11**, 687–712.

Biau, G. and Devroye, L. (2015). *Lectures on the Nearest Neighbor Method*. Springer, New York.

Boucheron, S., Bousquet, O. and Lugosi, G. (2005). Theory of classification: a survey of some recent advances. *ESAIM: PS*, **9**, 323–375.

Boucheron, S., Lugosi, G. and Massart, P. (2013). *Concentration Inequalities*. Oxford University Press, Oxford.

Celisse, A. and Mary-Huard, T. (2015). New upper bounds on cross-validation for the k-nearest neighbor classification rule. `https://arxiv.org/abs/1508.04905`.

Chapelle, O., Zien, A., and Schölkopf, B. (Eds.) (2006). *Semi-supervised Learning*. MIT Press, Cambridge MA.

Chaudhuri, K. and Dasgupta, S. (2014). Rates of convergence for nearest neighbor classification. *Advances in Neural Information Processing Systems*, **27**, 3437–3445.

Cover, T. M. and Hart, P. E. (1967). Nearest neighbour pattern classification. *IEEE Trans. Inf. Th.*, **13**, 21–27.

Devroye, L., Gyorfi, L. and Lugosi, G. (1996). *A Probabilistic Theory of Pattern Recognition*. Springer, New York.

Duda, R. O. and Hart, P. E. (1973). *Pattern Classification and Scene Analysis.* Wiley, New York.

Dudley, R. M. (1999). *Uniform Central Limit Theorems*. Cambridge University Press, Cambridge.

Duong, T. (2015). `ks`: Kernel smoothing. `R` package version 1.9.4, https://cran.r-project.org/web/packages/ks.

Fix, E. and Hodges, J. L. (1951). Discriminatory analysis – nonparametric discrimination: Consistency properties. Technical Report number 4, USAF School of Aviation Medicine, Randolph Field, Texas.

Fix, E. and Hodges, J. L. (1989). Discriminatory analysis – nonparametric discrimination: Consistency properties. *Internat. Statist. Rev.*, **57**, 238–247.

Gadat, S., Klein, T. and Marteau, C. (2016). Classification with the nearest neighbour rule in general finite dimensional spaces. *Ann. Statist.*, **44**, 982–1001.

Giné, E. and Guillou, A. (2002). Rates of strong uniform consistency for multivariate kernel density estimators. *Ann. Inst. H. Poincaré Probab. Statist.*, **38**, 907-921.

Gray, A. (2004). *Tubes, 2nd ed.* Progress in Mathematics 221. Birkhäuser, Basel.

Guillemin, V. and Pollack, A. (1974). *Differential Geometry*. Prentice-Hall, New Jersey.

Hall, P., and Kang, K.-H. (2005). Bandwidth choice for nonparametric classification. *Ann. Statist.*, **33**, 284–306.

Hall, P., Park, B. U. and Samworth, R. J. (2008). Choice of neighbour order in nearest-neighbour classification. *Ann. Statist.*, **36**, 2135–2152.

Heckel, R. and Bölsckei, H. (2015) Robust subspace clustering via thresholding. *IEEE Trans. Info. Th.*, **61**, 6320–6342.

Kozachenko, L. F. and Leonenko, N. N. (1987). Sample estimate of the entropy of a random vector. *Probl. Inform. Transm.*, **23**, 95–101.

Kulkarni, S. R. and Posner, S. E. (1995). Rates of convergence of nearest neighbor estimation under arbitrary sampling. *IEEE Trans. Info. Th.*, **41**, 1028–1039.

Mammen, E. and Tsybakov, A. B. (1999). Smooth discriminant analysis. *Ann. Statist.*, **27**, 1808–1829.

Paditz, L. (1989). On the analytical structure of the constant in the nonuniform version of the Esseen inequality. *Statistics*, **20**, 453–464.

Samworth, R. J. (2012). Optimal weighted nearest neighbour classifiers. *Ann. Statist.*, **40**, 2733–2763.

Schilling, M. F. (1986). Multivariate two-sample tests based on nearest neighbors. *J. Amer. Statist. Assoc.*, **81**, 799–806.

Shorack, G. R. and Wellner, J. A. (1986). *Empirical Processes with Applications to Statistics*. Wiley, New York.

Stone, C. J. (1977). Consistent nonparametric regression. *Ann. Statist.*, **5**, 595–620.

7 Appendix: An introduction to differential geometry, tubular neighbourhoods and integration on manifolds

The purpose of this section is to give a brief introduction to the ideas from differential geometry, specifically tubular neighbourhoods and integration on manifolds, which play an important role in our analysis of misclassification error rates, but which we expect are unfamiliar to many statisticians. For further details and several of the proofs, we refer the reader to the many excellent texts on these topics, e.g. Guillemin and Pollack (1974), Gray (2004).

7.1 Manifolds and regular values

Recall that if $\mathcal{X}$ is an arbitrary subset of $\mathbb{R}^M$, we say $\phi : \mathcal{X} \to \mathbb{R}^N$ is *differentiable* if for each $x \in \mathcal{X}$, there exists an open subset $U \subseteq \mathbb{R}^M$ containing x and a differentiable function $F : U \to \mathbb{R}^N$ such that $F(z) = \phi(z)$ for $z \in U \cap \mathcal{X}$. If $\mathcal{Y}$ is also a subset of $\mathbb{R}^M$, we say $\phi : \mathcal{X} \to \mathcal{Y}$ is a *diffeomorphism* if ϕ is bijective and differentiable and if its inverse ϕ^{-1} is also differentiable. We then say $\mathcal{S} \subseteq \mathbb{R}^d$ is an *m-dimensional manifold* if for each $x \in \mathcal{S}$, there exist an open subset $U_x \subseteq \mathbb{R}^m$, a neighbourhood V_x of x in $\mathcal{S}$ and a diffeomorphism $\phi_x : U_x \to V_x$. Such a diffeomorphism ϕ_x is called a *local parametrisation* of $\mathcal{S}$ around x, and we sometimes suppress the dependence of ϕ_x, U_x and V_x on x. It turns out that the specific choice of local parametrisation is usually not important, and properties of the manifold are well-defined regardless of the choice made.

Let $\mathcal{S} \subseteq \mathbb{R}^d$ be an m-dimensional manifold and let $\phi : U \to \mathcal{S}$ be a local parametrisation of $\mathcal{S}$ around $x \in \mathcal{S}$, where U is an open subset of $\mathbb{R}^m$. Assume that $\phi(0) = x$ for convenience.

The *tangent space* $T_x(\mathcal{S})$ to $\mathcal{S}$ at x is defined to be the image of the derivative $D\phi_0 : \mathbb{R}^m \to \mathbb{R}^d$ of ϕ at 0. Thus $T_x(\mathcal{S})$ is the m-dimensional subspace of $\mathbb{R}^d$ whose parallel translate $x + T_x(\mathcal{S})$ is the best affine approximation to $\mathcal{S}$ through x, and $(D\phi_0)^{-1}$ is well-defined as a map from $T_x(\mathcal{S})$ to $\mathbb{R}^m$. If $f : \mathcal{S} \to \mathbb{R}$ is differentiable, we define the derivative $Df_x : T_x(\mathcal{S}) \to \mathbb{R}$ of f at x by $Df_x := Dh_0 \circ (D\phi_0)^{-1}$, where $h := f \circ \phi$.

In practice, it is usually rather inefficient to define manifolds through explicit diffeomorphisms. Instead, we can often obtain them as level sets of differentiable functions. Suppose that $\mathcal{R} \subseteq \mathbb{R}^d$ is a manifold and $\eta : \mathcal{R} \to \mathbb{R}$ is differentiable. We say $y \in \mathbb{R}$ is a *regular value* for η if $\text{image}(D\eta_x) = \mathbb{R}$ for every $x \in \mathcal{R}$ for which $\eta(x) = y$. If $y \in \mathbb{R}$ is a regular value of η, then $\eta^{-1}(y)$ is a $(d-1)$-dimensional submanifold of $\mathcal{R}$ (Guillemin and Pollack, 1974, p. 21).

7.2 Tubular neighbourhoods of level sets

For any set $\mathcal{S} \subseteq \mathbb{R}^d$ and $\epsilon > 0$, we call $\mathcal{S} + \epsilon B_1(0)$ the *ϵ-neighbourhood* of $\mathcal{S}$. In circumstances where $\mathcal{S}$ is a $(d-1)$-dimensional manifold defined by the level set of a continuously differentiable function $\eta : \mathbb{R}^d \to \mathbb{R}$ with non-vanishing derivative on $\mathcal{S}$, the set $\mathcal{S}^\epsilon$ is often called a *tubular neighbourhood*, and $\dot\eta(x)^T v = 0$ for all $x \in \mathcal{S}$ and $v \in T_x(\mathcal{S})$. We therefore have the following useful representation of the ϵ-neighbourhood of $\mathcal{S}$ in terms of points on $\mathcal{S}$ and a perturbation in a normal direction.

Proposition 5. *Let $\eta : \mathbb{R}^d \to [0,1]$, suppose that $\mathcal{S} := \{x \in \mathbb{R}^d : \eta(x) = 1/2\}$ is non-empty, and suppose further that η is continuously differentiable on $\mathcal{S} + \epsilon B_1(0)$ for some $\epsilon > 0$, with $\dot\eta(x) \neq 0$ for all $x \in \mathcal{S}$, so that $\mathcal{S}$ is a $(d-1)$-dimensional manifold. Then*

$$\mathcal{S} + \epsilon B_1(0) = \left\{ x_0 + \frac{t\dot\eta(x_0)}{\|\dot\eta(x_0)\|} : x_0 \in \mathcal{S}, |t| < \epsilon \right\} =: \mathcal{S}^\epsilon.$$

Proof. For any $x_0 \in \mathcal{S}$ and $|t| < \epsilon$, we have $x_0 + t\dot\eta(x_0)/\|\dot\eta(x_0)\| \in \mathcal{S} + \epsilon B_1(0)$. On the other hand, suppose that $x \in \mathcal{S} + \epsilon B_1(0)$. Since $\mathcal{S}$ is closed, there exists $x_0 \in \mathcal{S}$ such that $\|x - x_0\| \leq \|x - y\|$ for all $y \in \mathcal{S}$. Rearranging this inequality yields that, for $y \neq x_0$,

$$2(x - x_0)^T \frac{(y - x_0)}{\|y - x_0\|} \leq \|y - x_0\|. \tag{46}$$

Let U be an open subset of $\mathbb{R}^{d-1}$ and $\phi : U \to \mathcal{S}$ be a local parametrisation of $\mathcal{S}$ around x_0, where without loss of generality we assume $\phi(0) = x_0$. Let $v \in T_{x_0}(\mathcal{S}) \setminus \{0\}$ be given and let $h \in \mathbb{R}^{d-1} \setminus \{0\}$ be such that $D\phi_0(h) = v$. Then for $t > 0$ sufficiently small we have $th \in U$, so by (46),

$$2(x - x_0)^T \frac{\{\phi(th) - \phi(0)\}}{\|\phi(th) - \phi(0)\|} \leq \|\phi(th) - \phi(0)\|.$$

Letting $t \searrow 0$ we see that $(x - x_0)^T v \leq 0$. Since $v \in T_{x_0}(\mathcal{S}) \setminus \{0\}$ was arbitrary and $-v \in T_{x_0}(\mathcal{S}) \setminus \{0\}$, we therefore have that $(x - x_0)^T v = 0$ for all $v \in T_{x_0}(\mathcal{S})$. Moreover, $\dot{\eta}(x_0)^T v = 0$ for all $v \in T_{x_0}(\mathcal{S})$, so $x - x_0 \propto \dot{\eta}(x_0)$, which yields the result. □

In fact, under a slightly stronger condition on η, we have the following useful result:

Proposition 6. *Let $\mathcal{R}$ be a d-dimensional manifold in $\mathbb{R}^d$, suppose that $\eta : \mathcal{R} \to [0, 1]$ satisfies the condition that $\mathcal{S} := \{x \in \mathcal{R} : \eta(x) = 1/2\}$ is non-empty. Suppose further that there exists $\epsilon > 0$ such that η is twice continuously differentiable on $\mathcal{S}^\epsilon$. Assume that $\dot{\eta}(x_0) \neq 0$ for all $x_0 \in \mathcal{S}$. Define $g : \mathcal{S} \times (-\epsilon, \epsilon) \to \mathcal{S}^\epsilon$ by*

$$g(x_0, t) := x_0 + \frac{t\dot{\eta}(x_0)}{\|\dot{\eta}(x_0)\|}.$$

If

$$\epsilon \leq \inf_{x_0 \in \mathcal{S}} \frac{\|\dot{\eta}(x_0)\|}{\sup_{z \in B_{2\epsilon}(x_0) \cap \mathcal{S}^\epsilon} \|\ddot{\eta}(z)\|_{\mathrm{op}}}, \tag{47}$$

then g is injective. In fact g is a diffeomorphism, with

$$Dg_{(x_0,t)}(v_1, v_2) = (I + tB)\left(v_1 + \frac{\dot{\eta}(x_0)}{\|\dot{\eta}(x_0)\|} v_2\right), \tag{48}$$

for $v_1 \in T_{x_0}(\mathcal{S})$ and $v_2 \in \mathbb{R}$, where

$$B := \frac{1}{\|\dot{\eta}(x_0)\|}\left(I - \frac{\dot{\eta}(x_0)\dot{\eta}(x_0)^T}{\|\dot{\eta}(x_0)\|^2}\right)\ddot{\eta}(x_0). \tag{49}$$

Proof. Assume for a contradiction that there exist distinct points $x_1, x_2 \in \mathcal{S}$ and $t_1, t_2 \in (-\epsilon, \epsilon)$ with $|t_1| \geq |t_2|$ such that

$$x_1 + \frac{t_1\dot{\eta}(x_1)}{\|\dot{\eta}(x_1)\|} = x_2 + \frac{t_2\dot{\eta}(x_2)}{\|\dot{\eta}(x_2)\|}.$$

Then

$$0 < \|x_2 - x_1\|^2 = \frac{2t_1\dot{\eta}(x_1)^T(x_2 - x_1)}{\|\dot{\eta}(x_1)\|} + t_2^2 - t_1^2 \leq \frac{2t_1\dot{\eta}(x_1)^T(x_2 - x_1)}{\|\dot{\eta}(x_1)\|}. \tag{50}$$

By Taylor's theorem and (50),

$$\begin{aligned}|\dot{\eta}(x_1)^T(x_2 - x_1)| = |\eta(x_2) - \eta(x_1) - \dot{\eta}(x_1)^T(x_2 - x_1)| &\leq \frac{1}{2} \sup_{z \in B_{2\epsilon}(x_1) \cap \mathcal{S}^\epsilon} \|\ddot{\eta}(z)\|_{\mathrm{op}} \|x_2 - x_1\|^2 \\ &< \sup_{z \in B_{2\epsilon}(x_1) \cap \mathcal{S}^\epsilon} \|\ddot{\eta}(z)\|_{\mathrm{op}} \frac{\epsilon|\dot{\eta}(x_1)^T(x_2 - x_1)|}{\|\dot{\eta}(x_1)\|},\end{aligned}$$

contradicting the hypothesis (47).

To show that g is a diffeomorphism, let $x_0 \in \mathcal{S}$ be given and let $\phi : U \to \mathcal{S}$ be a local parametrisation around x_0 with $\phi(0) = x_0$. Define $\Phi : U \times (-\epsilon, \epsilon) \to \mathcal{S} \times (-\epsilon, \epsilon)$ by $\Phi(u,t) := (\phi(u), t)$, and $H : U \times (-\epsilon, \epsilon) \to \mathcal{S}^\epsilon$ by $H := g \circ \Phi$. Finally, define the *Gauss map* $n : \mathcal{S} \to \mathbb{R}^d$ by $n(x_0) := \dot\eta(x_0)/\|\dot\eta(x_0)\|$. Then, for $h = (h_1^T, h_2)^T \in \mathbb{R}^{d-1} \times \mathbb{R}$ and $s \in \mathbb{R} \setminus \{0\}$,

$$\begin{aligned}
&\lim_{s\to 0} \frac{H(sh_1, t + sh_2) - H(0,t)}{s} \\
&\qquad = \lim_{s\to 0} \left\{ \frac{\phi(sh_1) - \phi(0)}{s} + \frac{t\{n(\phi(sh_1)) - n(\phi(0))\}}{s} + h_2 n\big(\phi(sh_1)\big) \right\} \\
&\qquad = D\phi_0(h_1) + tDn_{x_0} \circ D\phi_0(h_1) + h_2 n(x_0) \\
&\qquad = Dg_{(x_0,t)} \circ D\Phi_{(0,t)}(h_1, h_2),
\end{aligned}$$

where $Dg_{(x_0,t)} : T_{x_0}(\mathcal{S}) \times \mathbb{R} \to \mathbb{R}^d$ is given in (48).

To show that $Dg_{(x_0,t)}$ is invertible, note that for $v_1 \in T_{x_0}(\mathcal{S})$ and $|t| < \epsilon$,

$$\frac{|t|}{\|\dot\eta(x_0)\|} \left\| \left(I - \frac{\dot\eta(x_0)\dot\eta(x_0)^T}{\|\dot\eta(x_0)\|^2} \right) \ddot\eta(x_0) v_1 \right\| \le \frac{|t| \|\ddot\eta(x_0)\|_{\mathrm{op}}}{\|\dot\eta(x_0)\|} \|v_1\| < \|v_1\|,$$

where the final inequality follows from (47). Since $v_1 + \frac{t}{\|\dot\eta(x_0)\|}\left(I - \frac{\dot\eta(x_0)\dot\eta(x_0)^T}{\|\dot\eta(x_0)\|^2}\right)\ddot\eta(x_0)v_1$ and $n(x_0)v_2$ are orthogonal, it follows that $Dg_{(x_0,t)}$ is indeed invertible. The inverse function theorem (e.g. Guillemin and Pollack, 1974, p. 13) then gives that g is a local diffeomorphism, and moreover, by Guillemin and Pollack (1974, Exercise 5, p. 18) and the fact that g is bijective, we can conclude that g is in fact a diffeomorphism. □

7.3 Forms, pullbacks and integration on manifolds

Let V be a (real) vector space of dimension m. We say $T : V^p \to \mathbb{R}$ is a *p-tensor* on V if it is p-linear, and write $\mathcal{F}^p(V^*)$ for the set of p-tensors on V. If $T \in \mathcal{F}^p(V^*)$ and $S \in \mathcal{F}^q(V^*)$, we define their *tensor product* $T \otimes S \in \mathcal{F}^{p+q}(V^*)$ by

$$T \otimes S(v_1, \ldots, v_p, v_{p+1}, \ldots, v_{p+q}) := T(v_1, \ldots, v_p) S(v_{p+1}, \ldots, v_{p+q}).$$

Let S_p denote the set of permutations of $\{1, \ldots, p\}$. If $\pi \in S_p$ and $T \in \mathcal{F}^p(V^*)$, we can define $T^\pi \in \mathcal{F}^p(V^*)$ by $T^\pi(v) := T(v_{\pi(1)}, \ldots, v_{\pi(p)})$ for $v = (v_1, \ldots, v_p) \in V^p$. We say T is *alternating* if $T^\sigma = -T$ for all transpositions $\sigma : \{1, \ldots, p\} \to \{1, \ldots, p\}$. The set of alternating p-tensors on V, denoted $\Lambda^p(V^*)$, is a vector space of dimension $\binom{m}{p}$. The function $\mathrm{Alt} : \mathcal{F}^p(V^*) \to \Lambda^p(V^*)$ is defined by

$$\mathrm{Alt}(T) := \frac{1}{p!} \sum_{\pi \in S_p} (-1)^{\mathrm{sgn}(\pi)} T^\pi,$$

where $\mathrm{sgn}(\pi)$ denotes the sign of the permutation π. If $T \in \Lambda^p(V^*)$ and $S \in \Lambda^q(V^*)$, we define their *wedge product* $T \wedge S \in \Lambda^{p+q}(V^*)$ by

$$T \wedge S := \mathrm{Alt}(T \otimes S).$$

If W is another (real) vector space and $A : V \to W$ is a linear map, we define the *transpose* $A^* : \Lambda^p(W^*) \to \Lambda^p(V^*)$ of A by

$$A^*T(v_1, \ldots, v_p) := T(Av_1, \ldots, Av_p).$$

Let $\mathcal{S}$ be a manifold. A *p-form* ω on $\mathcal{S}$ is a function which assigns to each $x \in \mathcal{S}$ an element $\omega(x) \in \Lambda^p(T_x(\mathcal{S})^*)$. If ω is a p-form on $\mathcal{S}$ and θ is a q-form on $\mathcal{S}$, we can define their wedge product $\omega \wedge \theta$ by $(\omega \wedge \theta)(x) := \omega(x) \wedge \theta(x)$. For $j = 1, \ldots, m$, let $x_j : \mathbb{R}^m \to \mathbb{R}$ denote the coordinate function $x_j(y_1, \ldots, y_m) := y_j$. These functions induce 1-forms dx_j, given by $dx_j(x)(y_1, \ldots, y_m) = y_j$ (so $dx_j(x) = D(x_j)_x$ in our previous notation). Letting $\mathcal{I} := \{(i_1, \ldots, i_p) : 1 \le i_1 < \ldots < i_p \le m\}$, for $I = (i_1, \ldots, i_p) \in \mathcal{I}$, we write

$$dx_I := dx_{i_1} \wedge \ldots \wedge dx_{i_p}.$$

It turns out (Guillemin and Pollack, 1974, p. 163) that any p-form on an open subset U of $\mathbb{R}^m$ can be uniquely expressed as

$$\sum_{I \in \mathcal{I}} f_I \, dx_I, \tag{51}$$

where each f_I is a real-valued function on U.

Recall that the set of all ordered bases of a vector space V is partitioned into two equivalence classes, and an *orientation* of V is simply an assignment of a positive sign to one equivalence class and a negative sign to the other. If V and W are oriented vector spaces in the sense that an orientation has been specified for each of them, then an isomorphism $A : V \to W$ always either preserves orientation in the sense that for any ordered basis β of V, the ordered basis $A\beta$ has the same sign as β, or it reverses it. We say an m-dimensional manifold $\mathcal{X}$ is *orientable* if for every $x \in \mathcal{X}$, there exist an open subset U of $\mathbb{R}^m$, a neighbourhood V of x in $\mathcal{X}$ and a diffeomorphism $\phi : U \to V$ such that $D\phi_u : \mathbb{R}^m \to T_x(\mathcal{X})$ preserves orientation for every $u \in U$. A map like ϕ above whose derivative at every point preserves orientation is called an *orientation-preserving* map.

If $\mathcal{X}$ and $\mathcal{Y}$ are manifolds, ω is a p-form on $\mathcal{Y}$ and $\psi : \mathcal{X} \to \mathcal{Y}$ is differentiable, we define the *pullback* $\psi^*\omega$ of ω by ψ to be the p-form on $\mathcal{X}$ given by

$$\psi^*\omega(x) := (D\psi_x)^*\omega\big(\psi(x)\big).$$

If V is an p-dimensional vector space and $A : V \to V$ is linear, then $A^*T = (\det A)T$ for all $T \in \Lambda^p(V)$ (Guillemin and Pollack, 1974, p. 160).

If ω is an m-form on an open subset U of $\mathbb{R}^m$, then by (51), we can write $\omega = f\, dx_1 \wedge \ldots \wedge dx_m$. If ω is an integrable form on U (i.e. f is an integrable function on U), we can define the integral of ω over U by

$$\int_U \omega := \int_U f(x_1, \ldots, x_m)\, dx_1 \ldots dx_m,$$

where the integral on the right-hand side is a usual Lebesgue integral. Now let $\mathcal{S}$ be an m-dimensional orientable manifold that can be parametrised with a single chart, in the sense that there exists an open subset U of $\mathbb{R}^m$ and an orientation-preserving diffeomorphism $\phi : U \to \mathcal{S}$. Define the *support* of an m-form ω on $\mathcal{S}$ to be the closure of $\{x \in \mathcal{S} : \omega(x) \neq 0\}$. If ω is compactly supported, then its pullback $\phi^*\omega$ is a compactly supported m-form on U; moreover $\phi^*\omega$ is integrable, and we can define the integral over $\mathcal{S}$ of ω by

$$\int_{\mathcal{S}} \omega := \int_U \phi^*\omega. \tag{52}$$

Alternatively, we can suppose that ω is non-negative and measurable in the sense that $\phi^*\omega = f\, dx_1 \wedge \ldots \wedge dx_m$, say, with f non-negative and measurable on U. In this case, we can also define the integral of ω over $\mathcal{S}$ via (52).

More generally, integrals of forms over more complicated manifolds can be defined via partitions of unity. Recall (Guillemin and Pollack, 1974, p. 52) that if $\mathcal{X}$ is an arbitrary subset of $\mathbb{R}^M$, and $\{V_\alpha : \alpha \in A\}$ is a (relatively) open cover of $\mathcal{X}$, then there exists a sequence of real-valued, differentiable functions (ρ_n) on $\mathcal{X}$, called a *partition of unity* with respect to $\{V_\alpha : \alpha \in A\}$, with the following properties:

(i) $\rho_n(x) \in [0, 1]$ for all $n \in \mathbb{N}$;

(ii) Each $x \in \mathcal{X}$ has a neighbourhood on which all but finitely many functions ρ_n are identically zero;

(iii) Each ρ_n is identically zero except on some closed set contained in some V_α;

(iv) $\sum_{n=1}^{\infty} \rho_n(x) = 1$ for all $x \in \mathcal{X}$.

Now let $\mathcal{S} \subseteq \mathbb{R}^d$ be an m-dimensional, orientable manifold, so for each $x \in \mathcal{S}$, there exist an open subset U_x of $\mathbb{R}^m$, a neighbourhood V_x of x in $\mathcal{S}$ and an orientation-preserving diffeomorphism $\phi_x : U_x \to V_x$. If ω is a compactly supported m-form on $\mathcal{S}$ and (ρ_n) denotes

a partition of unity on $\mathcal{S}$ with respect to $\{V_x : x \in \mathcal{S}\}$, we can define the integral of ω over $\mathcal{S}$ by

$$\int_{\mathcal{S}} \omega := \sum_{n=1}^{\infty} \int_{\mathcal{S}} \rho_n \omega. \tag{53}$$

In fact, writing Ω for the compact support of ω, we can find a neighbourhood W_x of $x \in \Omega$, $x_1, \ldots, x_N \in \Omega$ and a finite subset N_j of $\mathbb{N}$ such that $\{\rho_n : n \notin N_j\}$ are identically zero on W_{x_j}, and such that

$$\int_{\mathcal{S}} \omega = \sum_{j=1}^{N} \sum_{n \in N_j} \int_{\mathcal{S}} \rho_n \omega.$$

Thus the integral can be written as a finite sum. Similarly, if ω is a non-negative m-form on $\mathcal{S}$, we can again define the integral of ω over $\mathcal{S}$ via (53). Finally, if ω is an integrable m-form on $\mathcal{S}$, the integral can be defined by taking positive and negative parts in the usual way.

In our work, we are especially interested in integrals of a particular type of form. Given an m-dimensional, orientable manifold $\mathcal{S}$ in $\mathbb{R}^d$, the *volume form* $d\mathrm{Vol}^m$ is the unique m-form on $\mathcal{S}$ such that at each $x \in \mathcal{S}$, the alternating m-tensor $d\mathrm{Vol}^m(x)$ on $T_x(\mathcal{S})$ gives value $1/m!$ to each positively oriented orthonormal basis for $T_x(\mathcal{S})$. For example, when $\mathcal{S} = \mathbb{R}^m$, we have $d\mathrm{Vol}^m = dx_1 \wedge \ldots \wedge dx_m$, provided we consider the standard basis to be positively oriented. As another example, if $\mathcal{R} \subseteq \mathbb{R}^d$ is a d-dimensional manifold and $\eta : \mathcal{R} \to \mathbb{R}$ is continuously differentiable with $\mathcal{S} = \{x \in \mathcal{R} : \eta(x) = 1/2\}$ non-empty and $\dot{\eta}(x) \neq 0$ for $x \in \mathcal{S}$, then $\mathcal{S}$ is a $(d-1)$-dimensional, orientable manifold (Guillemin and Pollack, 1974, Exercise 18, p. 106). If we say that an ordered, orthonormal basis $e_1, \ldots, e_{d-1}$ for $T_{x_0}(\mathcal{S})$ is positively oriented whenever $\det(e_1, \ldots, e_{d-1}, \dot{\eta}(x_0)) > 0$, we have that

$$d\mathrm{Vol}^{d-1}(x_0) = \sum_{j=1}^{d} (-1)^{j+d} \frac{\eta_j(x_0)}{\|\dot{\eta}(x_0)\|} dx_1 \wedge \ldots \wedge dx_{j-1} \wedge dx_{j+1} \wedge \ldots \wedge dx_d(x_0),$$

where x_j denotes the jth coordinate function. We now define an ordered, orthonormal basis $(e_1, 0), \ldots, (e_{d-1}, 0), (0, 1)$ for $T_{x_0}(\mathcal{S}) \times \mathbb{R}$ to be positively oriented. Further, we define a $(d-1)$-form ω_1 and a 1-form ω_2 on $\mathcal{S} \times (-\epsilon, \epsilon)$ by

$$\omega_1(x_0, t)\big((v_1, w_1), \ldots, (v_{d-1}, w_{d-1})\big) := d\mathrm{Vol}^{d-1}(x_0)(v_1, \ldots, v_{d-1})$$
$$\omega_2(x_0, t)(v_d, w_d) := dt(t)(w_d) = w_d.$$

Then, with g defined as in Proposition 6, and under the conditions of that proposition,

$$\begin{aligned}
&g^*(dx_1 \wedge \ldots \wedge dx_d)(x_0,t)\big((e_1,0),\ldots,(e_{d-1},0),(0,1)\big) \\
&\qquad = dx_1 \wedge \ldots \wedge dx_d(x_0^t)\big(Dg_{(x_0,t)}(e_1,0),\ldots,Dg_{(x_0,t)}(e_{d-1},0),Dg_{(x_0,t)}(0,1)\big) \\
&\qquad = \frac{1}{d!}\det(I+tB) \\
&\qquad = \frac{1}{d}\det(I+tB)d\mathrm{Vol}^{d-1}(x_0)(e_1,\ldots,e_{d-1})dt(t)(1) \\
&\qquad = \det(I+tB)\ (\omega_1 \wedge \omega_2)(x_0,t)\big((e_1,0),\ldots,(e_{d-1},0),(0,1)\big),
\end{aligned}$$

so $g^*(dx_1 \wedge \ldots \wedge dx_d)(x_0,t) = \det(I+tB)\ (\omega_1 \wedge \omega_2)(x_0,t)$. It follows that if $h : \mathcal{S} \times (-\epsilon,\epsilon) \to \mathbb{R}$ is either compactly supported and integrable, or non-negative and measurable, then

$$\int_{\mathcal{S}\times(-\epsilon,\epsilon)} h\,\omega_1 \wedge \omega_2 = \int_{\mathcal{S}} \int_{-\epsilon}^{\epsilon} h(x_0,t)\,dt\,d\mathrm{Vol}^{d-1}(x_0). \tag{54}$$

Finally, we require the change of variables formula: if $\mathcal{X}$ and $\mathcal{Y}$ are orientable manifolds and are of dimension m, and if $\psi : \mathcal{X} \to \mathcal{Y}$ is an orientation-preserving diffeomorphism, then

$$\int_{\mathcal{X}} \psi^*\omega = \int_{\mathcal{Y}} \omega \tag{55}$$

for every compactly supported, integrable m-form on $\mathcal{Y}$ (Guillemin and Pollack, 1974, p. 168). In particular, if $f : \mathcal{S}^\epsilon \to \mathbb{R}$ is either compactly supported and integrable, or non-negative and measurable, then writing $x_0^t := x_0 + \frac{t\dot{\eta}(x_0)}{\|\dot{\eta}(x_0)\|}$, we have from (54) and (55) that

$$\int_{\mathcal{S}^\epsilon} f(x)\,dx = \int_{\mathcal{S}\times(-\epsilon,\epsilon)} \det(I+tB) f(x_0^t)\ (\omega_1 \wedge \omega_2)(x_0,t) = \int_{\mathcal{S}} \int_{-\epsilon}^{\epsilon} \det(I+tB) f(x_0^t)\,dt\,d\mathrm{Vol}^{d-1}(x_0). \tag{56}$$